

Deep Predictive Learning in Neocortex and Pulvinar

Randall C. O'Reilly¹, Jacob L. Russin, Maryam Zolfaghar, and John Rohrlich

Abstract

■ How do humans learn from raw sensory experience? Throughout life, but most obviously in infancy, we learn without explicit instruction. We propose a detailed biological mechanism for the widely embraced idea that learning is driven by the differences between predictions and actual outcomes (i.e., predictive error-driven learning). Specifically, numerous weak projections into the pulvinar nucleus of the thalamus generate top-down predictions, and sparse driver inputs from lower areas supply the actual outcome, originating in Layer 5 intrinsic bursting neurons. Thus, the outcome representation is only briefly activated, roughly

every 100 msec (i.e., 10 Hz, alpha), resulting in a temporal difference error signal, which drives local synaptic changes throughout the neocortex. This results in a biologically plausible form of error backpropagation learning. We implemented these mechanisms in a large-scale model of the visual system and found that the simulated inferotemporal pathway learns to systematically categorize 3-D objects according to invariant shape properties, based solely on predictive learning from raw visual inputs. These categories match human judgments on the same stimuli and are consistent with neural representations in inferotemporal cortex in primates. ■

INTRODUCTION

The fundamental epistemological conundrum of how knowledge emerges from raw experience has challenged philosophers and scientists for centuries. Although there have been significant advances in cognitive and computational models of learning (LeCun, Bengio, & Hinton, 2015; Watanabe & Sasaki, 2015; Ashby & Maddox, 2011) and in our understanding of the detailed biochemical basis of synaptic plasticity (Cooper & Bear, 2012; Lüscher & Malenka, 2012; Urakubo, Honda, Froemke, & Kuroda, 2008; Shouval, Bear, & Cooper, 2002), there is still no widely accepted answer to this puzzle that is clearly supported by known biological mechanisms and also produces effective learning at the computational and cognitive levels. The idea that we learn via an active predictive process was advanced by Helmholtz in his “recognition by synthesis” proposal (von Helmholtz, 1867/2013) and has been widely embraced in a range of different frameworks (de Lange, Heilbron, & Kok, 2018; Summerfield & de Lange, 2014; Clark, 2013; George & Hawkins, 2009; Friston, 2005; Hawkins & Blakeslee, 2004; Rao & Ballard, 1999; Elman et al., 1996; Dayan, Hinton, Neal, & Zemel, 1995; Kawato, Hayakawa, & Inui, 1993; Mumford, 1992; Elman, 1990).

Here, we propose a detailed biological mechanism for a specific form of “predictive error-driven learning” based on distinctive patterns of connectivity between the neocortex and the higher-order nuclei of the thalamus (i.e., the pulvinar; Usrey & Sherman, 2018; Sherman & Guillery, 2006). We hypothesize that learning is driven by the difference between top-down predictions, generated by numerous weak

projections into the thalamic relay cells (TRCs) in the pulvinar, and the actual outcomes supplied by sparse, strong driver inputs from lower areas. Because these driver inputs originate in Layer 5 intrinsic bursting (5IB) neurons, the outcome is only briefly activated, roughly every 100 msec (i.e., 10 Hz, alpha). Thus, the prediction error is a temporal difference in activation states over the pulvinar, from an earlier prediction to a subsequent burst of outcome. This temporal difference can drive local synaptic changes throughout the neocortex, supporting a biologically plausible form of error backpropagation (Bp) that improves the predictions over time (Lillicrap, Santoro, Marris, Akerman, & Hinton, 2020; Whittington & Bogacz, 2019; Bengio, Mesnard, Fischer, Zhang, & Wu, 2017; O'Reilly, 1996; Hinton & McClelland, 1988; Ackley, Hinton, & Sejnowski, 1985). The temporal difference form of error-driven learning contrasts with prevalent alternative hypotheses that require a separate population of neurons to compute a prediction error explicitly and transmit it directly through neural firing (Lotter, Kreiman, & Cox, 2016; Ouden, Kok, & Lange, 2012; Friston, 2005, 2010; Rao & Ballard, 1999; Kawato et al., 1993).

In the following, our primary objective is to describe the hypothesized biologically based mechanism for predictive error-driven learning, contrast it with other existing proposals regarding the functions of this thalamocortical circuitry and other ways that the brain might support predictive learning, and evaluate it relative to a wide range of existing anatomical and electrophysiological data. We provide a number of specific empirical predictions that follow from this functional view of the thalamocortical circuit, which could potentially be tested by current neuroscientific methods. Thus, this work proposes a clear

¹University of California Davis

functional interpretation of this distinctive thalamocortical circuitry that contrasts with existing ideas in testable ways.

A second major objective is to implement this predictive error-driven learning mechanism in a large-scale computational model that faithfully captures its essential biological features, to test whether the proposed learning mechanism can drive the formation of cognitively useful representations. In particular, we ask a critical question for any predictive learning model: Can it develop high-level, abstract representations while learning from nothing but predicting low-level visual inputs? Most visual object recognition models that provide a reasonable fit to neurophysiological data rely on large human-labeled data sets to explicitly train abstract category information via error B (Rajalingham et al., 2018; Cadieu et al., 2014; Khaligh-Razavi & Kriegeskorte, 2014). Thus, it is perhaps not too surprising that the higher layers of these models, which are closer to these category output labels, exhibited a greater degree of categorical organization.

Through large-scale simulations based on the known structure of the visual system, we found that our biologically based predictive learning mechanism developed high-level, abstract representations that significantly diverge from the similarity structure present in the lower layers of the network and systematically categorize 3-D objects according to invariant shape properties. Furthermore, we found in an experiment using the same stimuli that these categories match human similarity judgments and that they are also qualitatively consistent with neural representations in inferotemporal (IT) cortex in primates (Cadieu et al., 2014). In addition, we show that comparison predictive Bp models lacking these biological features (Lotter et al., 2016) did not learn object categories that go beyond the visual input structure. Thus, there may be some important features of the biologically based model that enable this ability to learn higher-level structure beyond that of the raw inputs.

It is important to emphasize that our objectives for these simulations are not to produce a better machine-learning algorithm per se but rather to test whether our biologically


based model can capture some of the known high-level, cognitive phenomena that the mammalian brain learns. Thus, we explicitly dissuade readers from the inevitable desire to evaluate the importance of our model based on differences in narrow, performance-based machine learning metrics. As discussed later, there are various engineering-level issues regarding the biologically based model's computational cost and performance, which currently limit its ability to compete with simpler, much larger-scale Bp models, but we do not think these are relevant to the evaluation of the scientific questions of relevance here. In short, this model is an instantiation of a scientific theory, and it should be evaluated on its ability to explain a wide range of data across multiple levels of analysis, just as every other scientific theory is evaluated.

The remainder of the paper is organized as follows. First, we provide a concise overview of the biologically based predictive error-driven learning framework, including the most relevant neural data. Then, we present a small-scale implementation of the model that learns a probabilistic grammar, to illustrate the basic computational mechanisms of the theory. This is followed by the large-scale model of the visual system, which learns by predicting over brief movies of 3-D objects rotating and translating in space. We evaluate this model and compare it to two other predictive learning models that directly use error Bp, based on current deep convolutional neural network (DCNN) mechanisms. Then, we circle back to discuss the relevant biological data in greater detail, along with testable predictions that can differentiate this account from other existing ideas. Finally, we conclude with a discussion of related models and outstanding issues.

PREDICTIVE ERROR-DRIVEN LEARNING IN THE NEOCORTEX AND PULVINAR

Figure 1 shows the thalamocortical circuits characterized by Sherman and Guillery (2006; see also Usrey & Sherman, 2018; Sherman & Guillery, 2013), which have two distinct projections converging on the principal TRCs of the pulvinar,

Figure 1. Summary figure from Sherman and Guillery (2006) showing the strong feedforward driver projection emanating from layer 5IB cells in lower layers (e.g., VI) and the much more numerous feedback “modulatory” projection from layer 6CT (corticothalamic) cells. We interpret these same connections as providing a prediction (6CT) versus outcome (5IB) activity pattern over the pulvinar.



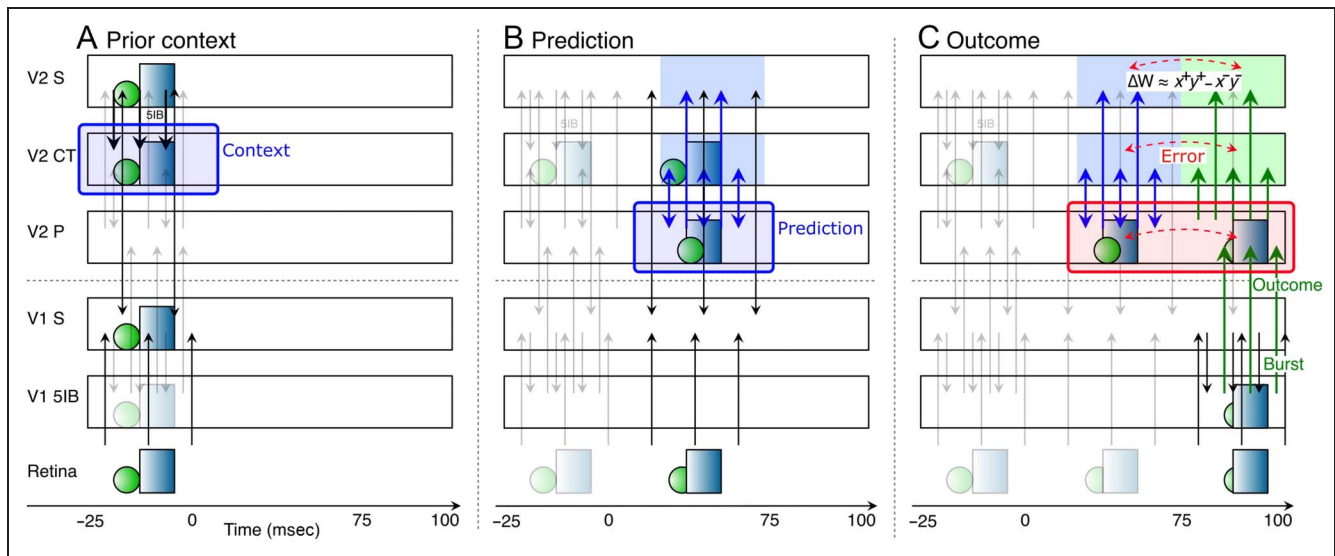


Figure 2. Corticothalamic (CT) information flow under our predictive learning hypothesis, shown as a sequence of movie frames (“Retina”), illustrating the three key steps taking place within a single 125-msec time window, broken out separately across the three panels: (A) Prior context is updated in the V2 CT layer, (B) which is then used to generate a prediction over the pulvina (V2 P), (C) against which the outcome, driven by bottom-up 5IB, represents the prediction error as a temporal difference between the prediction and outcome states over the pulvina. Changes in synaptic weights (learning) in all superficial (S) and CT layers are driven from the local temporal difference experienced by each neuron, using a form of the CHL term as shown, where the “+”: superscripts indicate outcome activations and “-” superscripts indicate prediction. CHL approximates the backpropagated prediction error gradient experienced by each neuron (O’Reilly, 1996), reflecting both direct pulvina error signals and indirect corticocortical error signals as well. In specific, (A) CT context updating occurs via 5IB (not shown) in higher layer (V2) during prior alpha (100-msec) cycle—this context is maintained in the CT layer and used to generate predictions. (B) The prediction state also projects up to S and CT layers, and from S to all other S layers via extensive bidirectional connectivity, so their activation state reflects this prediction as well. (C) The subsequent outcome drives pulvina activity bottom-up via V1 5IB and is likewise projected to S and CT layers, ensuring that the relevant temporal difference error signal is available locally in the cortex. The difference in activation values across these two time points, in S and CT layers throughout the network, drives learning to reduce prediction errors. Note that the single most important property of the 5IB is that these driver cells are not active during the prediction phase—the bursting itself may also be useful in the driving property, but that is a secondary consideration to the critical feature of having a time when the prediction alone can be projected onto the pulvina.

the primary thalamic nucleus that is interconnected with higher-level posterior cortical visual areas (Halassa & Kastner, 2017; Arcaro, Pinsk, & Kastner, 2015; Shipp, 2003). One projection consists of numerous, weaker connections originating in deep layer VI of the neocortex (the 6CT corticothalamic projecting cells), which we hypothesize generate a top-down prediction on the pulvina. The other is a sparse (Rockland, 1996, 1998) and strong driver pathway that originates from lower-level layer 5IB cells, which we hypothesize provide the outcome. These 5IB neurons fire discrete bursts with intrinsic dynamics having a period of roughly 100 msec between bursts (Saalmann, Pinsk, Wang, Li, & Kastner, 2012; Larkum, Zhu, & Sakmann, 1999; Franceschetti et al., 1995; Silva, Amitai, & Connors, 1991; Connors, Gutnick, & Prince, 1982), which is thought to drive the widely studied alpha frequency of ~10 Hz that originates in cortical deep layers and has important effects on a wide range of perceptual and attentional tasks (Clayton, Yeung, & Kadosh, 2018; Jensen, Bonnefond, & VanRullen, 2012; Buffalo, Fries, Landman, Buschman, & Desimone, 2011; Mathewson, Gratton, Fabiani, Beck, & Ro, 2009; VanRullen & Koch, 2003). Critically, unlike many other such bursting phenomena, this 5IB occurs in awake animals (Luczak, Bartho, and Harris, 2009, 2013; Sakata & Harris,

2009, 2012), consistent with the presence of alpha in awake, behaving states.

The existing literature generally characterizes the 6CT projection as modulatory (Usrey & Sherman, 2018; Sherman & Guillery, 2013), but a number of electrophysiological recordings from awake, behaving animals clearly show sustained, continuous patterns of neural firing in pulvina TRC neurons, which is not consistent with the idea that they are only being driven by their phasic bursting 5IB inputs (Zhou, Schafer, & Desimone, 2016; Komura, Nikkuni, Hirashima, Uetake, & Miyamoto, 2013; Saalmann et al., 2012; Bender & Youakim, 2001; Robinson, 1993; Petersen, Robinson, & Keys, 1985; Bender, 1982). Indeed, these recordings show that pulvina neural firing generally resembles that of the visual areas with which they interconnect, in terms of neural receptive field properties, tuning curves, and so forth. This is important because our predictive learning framework requires that these 6CT top-down projections be capable of directly driving TRC activity. Specifically, in contrast to the standard view, the core idea behind our theory is that the top-down 6CT projections drive a predicted activity pattern across the extent of the pulvina, which precedes the subsequent outcome activation state driven by the strong 5IB inputs.

Figure 2 illustrates the temporal evolution of activity states according to our predictive learning theory, which is somewhat challenging to convey because the critical signals driving learning unfold over time (O'Reilly, Wyatte, & Rohrlich, 2014, 2017; Kachergis, Wyatte, O'Reilly, de Kleijn, & Hommel, 2014). We hypothesize that synaptic plasticity throughout the cortex is sensitive to the resulting temporal differences that emerge initially in the pulvinar. Thus, unlike other models (as we discuss in depth later), the prediction error here is not captured directly in the firing of a special population of error-coding neurons but rather remains as a temporal difference error signal.

Figure 2 shows a single 125-msec time window of a 100-msec alpha cycle for the purposes of illustration (the actual timing is likely to be more dynamic as discussed next). The activity state in pulvinar TRC neurons, representing a prediction, as driven by the top-down 6CT projections, should develop during the first ~75 msec, when the 5IB neurons are paused between bursting. Then, the final ~25 msec largely reflects the strong 5IB bottom-up ground-truth driver inputs when they burst. Thus, the prediction error signal is reflected in the temporal difference of these activation states as they develop over time. In other words, our hypothesis is that the pulvinar is directly representing either the top-down prediction or the bottom-up outcome at any given time, and the temporal difference between these states implicitly encodes a prediction error. Whereas the deep 6CT layer is involved in generating a top-down prediction over the pulvinar, the superficial layer neurons continuously represent the current state, simultaneously incorporating bottom-up and top-down constraints via their own connections with other areas. To ensure that the prediction is not directly influenced by this current state representation (i.e., “peeking at the right answer”), it is important that the 6CT neurons encode temporally delayed information, consistent with available data (Harris & Shepherd, 2015; Thomson, 2010; Sakata & Harris, 2009).

The actual biological system is likely to be much more dynamic than the simplistic cartoon with rigid 100-msec timing, as shown in Figure 2, based on a set of neural mechanisms that can work together to enable it to more flexibly entrain the predictive learning cycle to the environment. These mechanisms would also tend to increase activity and learning associated with unexpected outcomes relative to expected ones, consistent with the observed expectation suppression phenomena (Bastos et al., 2012; Meyer & Olson, 2011; Todorovic, van Ede, Maris, & de Lange, 2011; Summerfield, Trittschuh, Monti, Mesulam, & Egnér, 2008).

Specifically, various underlying mechanisms result in neural adaptation, which is generally thought to increase neural activity and learning associated with novel inputs relative to recently familiar ones (Hennig, 2013; Grill-Spector, Henson, & Martin, 2006; Brette & Gerstner, 2005; Müller, Metha, Krauskopf, & Lennie, 1999; Abbott, Varela, Sen, & Nelson, 1997). In the case where outcomes are consistent with prior predictions (i.e., the predictions are accurate),

the same population of neurons across pulvinar and cortex should be active over time, whereas unpredicted outcomes will generally activate new subsets of neurons in superficial cortical layers representing the current state. Thus, because of adaptation, there should be a phasic increase in activity in these superficial neurons at the onset of unpredicted stimuli relative to predicted ones. Furthermore, the 5IB neurons downstream of these superficial neurons may be particularly responsive to these phasic activity increases, causing their bursting to coincide preferentially with unexpected outcomes, thereby driving the phase resetting of the alpha cycle to such events. Thus, during a sequence of predicted states, the pulvinar may experience relatively weaker or even absent 5IB driving inputs, until an unpredicted stimulus arises. At this point, error-driven learning would be more strongly engaged as a function of the phasic release from adaptation and 5IB burst activation. We discuss these dynamics more later in the context of the comparison with explicit error (EE) coding models.

We also hypothesize that 5IB preferentially drives the synaptic plasticity processes to take place at that time, because of the strong driving nature of the outputs from these neurons. In computational terms originating with the Boltzmann machine (Hinton & Salakhutdinov, 2006; Ackley et al., 1985), this anchors the target or plus phase to be at this point of 5IB. Furthermore, this means that the predictive nature of the prior minus phase naturally emerges just by virtue of it being the state before 5IB: The learning rule automatically causes that prior state to better anticipate the subsequent state. Thus, even if no prediction was initially generated, learning over multiple iterations will work to create one, to the extent that a reliable prediction can be generated based on internal states and environmental inputs. Likewise, assuming relevant activity traces naturally persist over timescales longer than the alpha cycle, this predictive learning process can take advantage of any such remaining traces to learn across these longer timescales, although it is operating at the faster alpha scale.

In short, learning always happens whenever something unexpected occurs, at any point, and drives the development of predictions immediately prior, to the extent such predictions are possible to generate. In the typical laboratory experiment where phasic stimuli are presented without any predictable temporal sequence (which is uncharacteristic of the natural world), there may often be no significant prediction before stimulus onset, and we would expect such stimuli to reliably drive 5IB, which is consistent with available electrophysiological data (Zhou et al., 2016; Komura et al., 2013; Luczak et al., 2009, 2013; Bender & Youakim, 2001; Robinson, 1993; Petersen et al., 1985; Bender, 1982). Thus, unlike Figure 2, such situations would start with a 5IB-triggered plus phase, without a significant minus phase before that.

As may be evident by this point, we are mainly focused on prediction in the sense of the humorous quote: “Prediction is very difficult, especially about the future” (attributable

to Danish author Robert Storm Petersen), whereas this term is potentially confusingly used in a much broader sense in most Bayesian-inspired predictive coding frameworks (de Lange et al., 2018; Friston, 2005; Rao & Ballard, 1999). These frameworks use “prediction” to encompass everything from genetic biases to the results of learning in the feedforward synaptic pathways to top-down filling-in or biasing of the current stimulus properties and fairly rarely use it in the “about the future” sense. We think these different phenomena are each associated with different neural mechanisms at different timescales (O’Reilly, Hazy, & Herd, 2016; O’Reilly, Wyatte, Herd, Mingus, & Jilk, 2013; O’Reilly, Munakata, Frank, Hazy, & Contributors, 2012) and thus prefer to treat them separately, while also recognizing that they can clearly interact as well.

Thus, our use of the term “prediction” here refers specifically to “anticipatory” neural firing that predicts subsequent stimuli. We use the term “postdiction” to refer to the operation of this predictive mechanism after a stimulus has been initially processed (to consolidate and more deeply encode, as in an autoencoder model) and distinguish both from top-down excitatory biasing, which directly influences the online superficial layer neural representations of the current stimulus (O’Reilly et al., 2013; Miller & Cohen, 2001; Reynolds, Chelazzi, & Desimone, 1999; Desimone & Duncan, 1995). Finally, many discussions of prediction error in the literature include late, frontally associated processes such as those associated with the P300 ERP component (Holroyd & Coles, 2002). We specifically exclude these from the scope of the mechanisms described here, which are anticipatory, fast, and low level, as is appropriate for the posterior cortical sensory processing areas that interconnect with the pulvinar.

Computational Properties of Predictive Learning in the Thalamocortical Circuits

We next elaborate the connections between the computational properties required for predictive learning and the properties of the circuits interconnecting the cortex and the pulvinar, which appear to be notably well suited for their hypothesized role in predictive learning. We begin with a relatively established interpretation of superficial layer processing, to contextualize subsequent points about the special functions required of the deep layers and the thalamus.

- The superficial cortical layers continuously represent the current state: The superficial layer pyramidal neurons are densely and bidirectionally interconnected with other cortical areas and update quickly to new stimulus inputs, with continuous, relatively rapid firing (i.e., up to about 100 Hz for preferred stimuli). These neurons integrate higher-level top-down information with bottom-up sensory information to resolve ambiguities, focus attention, fill in missing information, and generally enhance the consistency and quality of the

online representations (O’Reilly, Hazy, & Herd, 2016; O’Reilly, Wyatte, Herd, Mingus, & Jilk, 2013; O’Reilly et al., 2012; Miller & Cohen, 2001; Reynolds et al., 1999; Desimone & Duncan, 1995; Hopfield, 1984; Rumelhart & McClelland, 1982). As noted above, we distinguish this form of top-down processing, which is often most evident during the period after stimulus onset (Lee & Mumford, 2003), from the specifically predictive, anticipatory sort.

- Predictions must be insulated against receiving current state information (it is not prediction if you already know what happens): Given that the superficial layers are continuously updating and representing the current state, some kind of separate neural system insulated from this current state information must be used to generate predictions; otherwise, the prediction system can just “cheat” and directly report the current state. It may seem counterintuitive, but making the prediction task harder is actually beneficial, because that pushes the learning to capture deeper, more systematic regularities about how the environment evolves over time. In other words, like any kind of cheating, the cheater itself is cheated because of the reduced pressure to learn, and learning is the real goal.
- Predictions take time and space to generate: Nontrivial predictions likely require the integration of multiple converging inputs from a range of higher-level cortical areas, each encoding different dimensions of relevance (e.g., location, motion, color, texture, shape). Thus, sufficient time and space (i.e., neural substrates with relevant connectivity) must be available to integrate these signals into a coherent predicted state, and per the above point, these substrates must be separated from the influence of current state information. This fits with the properties of the layer 6CT neurons and their deep layer inputs, which we hypothesize are insulated from superficial-layer firing by virtue of being driven locally by the 5IB within their own cortical microcolumn, such that the interbursting pause period provides a time window when these deep layers can integrate and generate the prediction.
- Biologically, this is consistent with the delayed responses of 6CT neurons (Harris & Shepherd, 2015; Thomson, 2010; Sakata & Harris, 2009). Computationally, these neurons function much like the simple recurrent network (SRN) context layer updating (Elman, 1990; Jordan, 1989), which reflects the prior trial’s state, as discussed in detail in the Appendix. The overall duration of the alpha cycle may represent a reasonable compromise between the prediction integration time and the need to keep up with predictions tracking changes in the world. Notably, films are typically shown at only over two times the alpha frequency (24 Hz), suggesting a Nyquist sampling relative to the underlying alpha processing.
- The predicted state must be directly aligned with the outcome state it predicts: A prediction error is a difference between two states, so these prediction and

outcome states must be directly comparable such that their difference meaningfully represents the actual prediction error and not some other kind of irrelevant encoding differences. In other words, the prediction and the outcome must be represented in the same “language,” so that the “words” from the prediction can be directly compared against those of the outcome—if the prediction was in Japanese and the outcome was in English, it would be hard to tell whether the prediction was correct or not. Thus, a common neural substrate with two different input pathways is required, one reflecting the prediction and the other reflecting the outcome, so that both converge onto the same representational system within this common neural substrate. This fits well with the two pathways converging into the pulvinar: the 6CT top-down prediction-generation pathway and the lower-level 5IB driving inputs.

- The outcome signal should be as veridical as possible (i.e., directly reflecting the bottom-up outcome) and should arise from lower areas in the hierarchy relative to the corresponding predictive 6CT inputs: Given that the outcome is the driver of learning, if it were to be corrupted or inaccurate, then everything that is learned would then be suspect. To the extent that delusional thinking is present in all people (some more so than others perhaps), this principle must be violated at some level, but for the lowest levels of the perceptual system at least, it is important that strongly grounded, accurate training signals drive learning. The bottom-up, sparse, strongly driving nature of the 5IB projections to the pulvinar can directly convey such veridical outcome signals and ensure that they dominate the activation of their TRC targets. On the basis of indirect available data, it is likely that each pulvinar TRC neuron receives only roughly one to six driver inputs (Sherman & Guillery, 2006, 2011), such that these sparse inputs directly convey the signal from lower layers, without much further mixing or integration (which could distort the nature of the signal). Furthermore, these inputs are likely not plastic (Usrey & Sherman, 2018), again consistent with a need for unaltered, veridical signals. Finally, the TRC neurons are distinctive in having no significant lateral interconnectivity (Sherman & Guillery, 2006), enabling them to faithfully represent their inputs. These properties led Mumford (1991) to characterize the pulvinar as a blackboard, and we further suggest the metaphor of a projection screen upon which the predictions are projected.
- The prediction error must drive learning to reduce subsequent prediction errors: Obviously, this is the goal of prediction error learning in the first place, and given that the cortex is what generates predictions, it must be capable of learning based on prediction error signals represented over the pulvinar. Computationally, the critical problem here is “credit assignment”: How do the error signals direct learning in the proper direction for each individual neuron, to reduce the overall

prediction error? The error Bp procedure solves this problem (Rumelhart, Hinton, & Williams, 1986) but requires biologically implausible retrograde signaling across the entire network of neural communication (Crick, 1989), to propagate the error proportionally back along the same channels that drive forward activation. Bidirectional connections, which are ubiquitous in the cortex (Markov, Ercsey-Ravasz, et al., 2014; Felleman & Van Essen, 1991) and computationally beneficial for other reasons as noted earlier, can eliminate that problem by “implicitly” propagating error signals via standard neural communication mechanisms along both directions of connectivity (O’Reilly, 1996).

- This solution to the credit assignment problem relies on a temporal difference error signal, as originally developed for the Boltzmann machine (Ackley et al., 1985). The bidirectional neural communication at one point in time is encoding and sharing the prediction among the entire network of neurons. Then, this same network of connections is reused at another point in time to encode and communicate the outcome. Mathematically, the difference in activation state across these two points in time, locally at each individual neuron, provides an accurate estimate of the error Bp gradient (O’Reilly, 1996). In effect, this temporal difference tells each neuron which direction it needs to change its activation state to reduce the overall error. The reuse of the very same network of connections across both points in time ensures the overall alignment of the two activation states, as noted above, such that this temporal difference precisely represents the error signal. Although various other schemes for error-driven learning in biologically plausible networks have been proposed (e.g., Lillicrap et al., 2020; Whittington & Bogacz, 2019; Bengio et al., 2017), the temporal difference framework with bidirectional connectivity provides a particularly good fit with the natural temporal ordering of predictive learning (prediction and then outcome) and the extensive bidirectional connectivity of the thalamocortical circuits (Shipp, 2003).
- Temporal differences in activation state across the alpha cycle, between prediction and outcome states, must drive synaptic plasticity: The final step needed to connect all of the elements above is that neurons actually modify their synaptic strengths in proportion to the temporal difference error signal. We have recently provided a fully explicit mechanism for this form of learning (O’Reilly et al., 2012), based on a biologically detailed model of spike-timing-dependent plasticity (Urakubo et al., 2008). We showed that, when activated by realistic Poisson spike trains, this spike-timing-dependent plasticity model produces a nonmonotonic learning curve similar to that of the Bienenstock, Cooper, and Munro (BCM) model (Bienenstock, Cooper, & Munro, 1982), which results from competing calcium-driven postsynaptic plasticity pathways (Cooper & Bear, 2012; Shouval et al., 2002). As in the BCM framework, we hypothesized

that the threshold crossover point in this nonmonotonic curve moves dynamically—if this happens on the alpha timescale (Lim et al., 2015), then it can reflect the prediction phase of activity, producing a net error-driven learning rule based on a subsequent calcium signal reflecting the outcome state. The resulting learning mechanism naturally supports a combination of both BCM-style Hebbian learning and error-driven learning, where the BCM component acts as a kind of regularizer or bias, similar to weight decay (O'Reilly et al., 2012; O'Reilly & Munakata, 2000).

Thus, remarkably, the pulvinar and associated thalamocortical circuitry appear to provide precisely the necessary ingredients to support predictive error-driven learning, according to the above analysis. Interestingly, although Sherman and Guillery (2006) did not propose a predictive learning mechanism as just described, they did speculate about a potential role for this circuit in motor forward-model learning and the predictive remapping phenomenon (Usrey & Sherman, 2018; Sherman & Guillery, 2011). In addition, Pennartz, Dora, Muckli, and Lorteije (2019) also suggested that the pulvinar may be involved in predictive learning, but within the EE coding framework and not involving the detailed aspects of the above-described circuitry.

It bears emphasizing the synergy between the various considerations above for the benefits of the pause in 5IB firing between bursts. First, this pause is critical for creating the time window when the predictive network is representing and communicating the prediction state, without influence from the outcome state. Furthermore, it creates the temporal difference in activation state in the pulvinar between prediction and outcome, which is needed for driving error-driven learning. Thus, for both the 6CT and pulvinar layers, the periodic pausing of 5IB neurons is essential for creating the predictive learning dynamic. Interestingly, by these principles, the lack of such burst/pause dynamics in the driver inputs to first-order sensory thalamus areas such as the lateral geniculate nucleus and medial geniculate nucleus (Sherman & Guillery, 2006) means that these areas should not be directly capable of error-driven predictive learning. This is consistent with a number of models and theoretical proposals suggesting that primary sensory areas may learn predominantly through Hebbian-style self-organizing mechanisms (Bednar, 2012; Miller, 1994). Nevertheless, primary sensory areas do receive “collateral” error signals from the pulvinar (Shipp, 2003), which could provide some useful indirect error-driven learning signals.

Note that this form of temporal difference learning signal is distinct from the widely used temporal-difference model in reinforcement learning (Sutton & Barto, 1998), which is scalar and applies to reward expectations, not sensory predictions (although see Gardner, Schoenbaum, & Gershman, 2018, and Dayan, 1993, for potential connections between these two forms of prediction error). Finally, as we discuss later, this proposed predictive role for the pulvinar is

compatible with the more widely discussed role it may play in attention (Fiebelkorn & Kastner, 2019; Zhou et al., 2016; Saalman & Kastner, 2011; Snow, Allen, Rafal, & Humphreys, 2009; Bender & Youakim, 2001; LaBerge & Buchsbaum, 1990). Indeed, we think these two functions are synergistic (i.e., you predict what you attend, and vice versa; Richter & de Lange, 2019) and have initial computational results consistent with this idea.

PREDICTIVE LEARNING OF TEMPORAL STRUCTURE IN A PROBABILISTIC GRAMMAR

To illustrate and test the predictive learning abilities of this biologically based model, we first ran a classical test of sequence learning (Cleeremans & McClelland, 1991; Reber, 1967) that has been explored using SRNs (Elman, 1990; Jordan, 1989). The biologically based model was implemented using the Leabra algorithm, which is a comprehensive framework that uses conductance-based point neuron equations, inhibitory competition, bidirectional connectivity, and the biologically plausible temporal difference learning mechanism described above (O'Reilly et al., 2012, 2016; O'Reilly & Munakata, 2000; O'Reilly, 1996, 1998). Leabra serves as a model of the bidirectionally connected processing in the cortical superficial layers and has been used to simulate a large number of different cognitive neuroscience phenomena. It is described in the Appendix, which also provides a detailed mapping between the SRN and our biological model.

As shown in Figure 3, sequences were generated according to a finite state automaton (FSA) grammar, as used in implicit sequence learning experiments by Reber (1967). Each node has a 50% random branching to two different other nodes, and the labels generated by node transitions are locally ambiguous (except for the B = begin and E = end states). Thus, integration over time and across many iterations is required to infer the systematic underlying grammar. It is a reasonably challenging task for SRNs and people to learn and provides an important validation of the power of these predictive learning mechanisms. Given the random branching, accurately predicting the specific path taken is impossible, but we can score the model's output as correct if it activates either or both of the possible branches for each state.

The model (Figure 4) required around 20 epochs of 25 sequences through the grammar to learn it to the point of making no prediction errors for five epochs in a row (which guarantees that it had completely learned the task). This model is available in the standard emergent distribution at github.com/emergent/leabra/tree/master/examples/deep_fsa. A few steps through a sequence are shown in Figure 4, illustrating how the corticothalamic (CT) context layer, which drives the P pulvinar layer prediction, represents the information present on the previous alpha cycle time step. Thus, the network is attempting to predict the current input state, which then drives the pulvinar plus phase at the end of each alpha cycle, as




Figure 3. FSA grammar used in implicit sequential learning experiments (Reber, 1967) and in early SRNs (Cleeremans & McClelland, 1991). It generates a sequence of letters according to the link transitioned between state nodes, where each outgoing link to another node has a 50% probability of being selected. Each letter (except for the B = begin and E = end) appears at two different points in the grammar, making them locally ambiguous. This combination of randomness and ambiguity makes it challenging for a learning system to infer the true underlying structure of the grammar.

shown in the last panel. On each trial, the difference between plus and minus phases locally over each cortical neuron drives its synaptic weight changes, which accumulate over trials to allow accurate prediction of the sequences, to the extent possible given their probabilistic nature.

PREDICTIVE LEARNING OF OBJECT CATEGORIES IN IT CORTEX

Now, we describe a large-scale, systems-neuroscience implementation of the proposed thalamocortical predictive error-driven learning framework, in a model of visual predictive learning (Figure 5). Our second major objective, and a critical question for predictive learning, is determining whether the model can develop high-level, abstract ways of representing the raw sensory inputs, while learning from nothing but predicting these low-level visual inputs. We showed the model brief movies of 156 3-D object exemplars drawn from 20 different basic-level categories (e.g., car, stapler, table lamp, traffic cone) selected for their overall shape diversity from the CU3D-100 data set (O'Reilly et al., 2013). The objects moved and rotated in 3-D space over eight movie frames, where each frame was sampled at the alpha frequency (Figure 5B). Because the motion and rotation parameters were generated at random on each sequence, this data set consists of 512,000 unique images, and there is no low-dimensional object category training signal, so the usual concerns about overfitting and training versus testing sets are not applicable: Our main question is what kind of representations self-organize as a result of this purely visual experience.

There were also saccadic eye movements every other frame, introducing an additional, realistic, predictive learning challenge. An efferent copy signal enabled full prediction of the effects of the eye movement and allows the model to capture the signature predictive remapping phenomenon

(Neupane, Guitton, & Pack, 2017; Cavanagh, Hunt, Afraz, & Rolfs, 2010; Duhamel, Colby, & Goldberg, 1992). The only learning signal available to the model was the prediction error generated by the temporal difference between what it


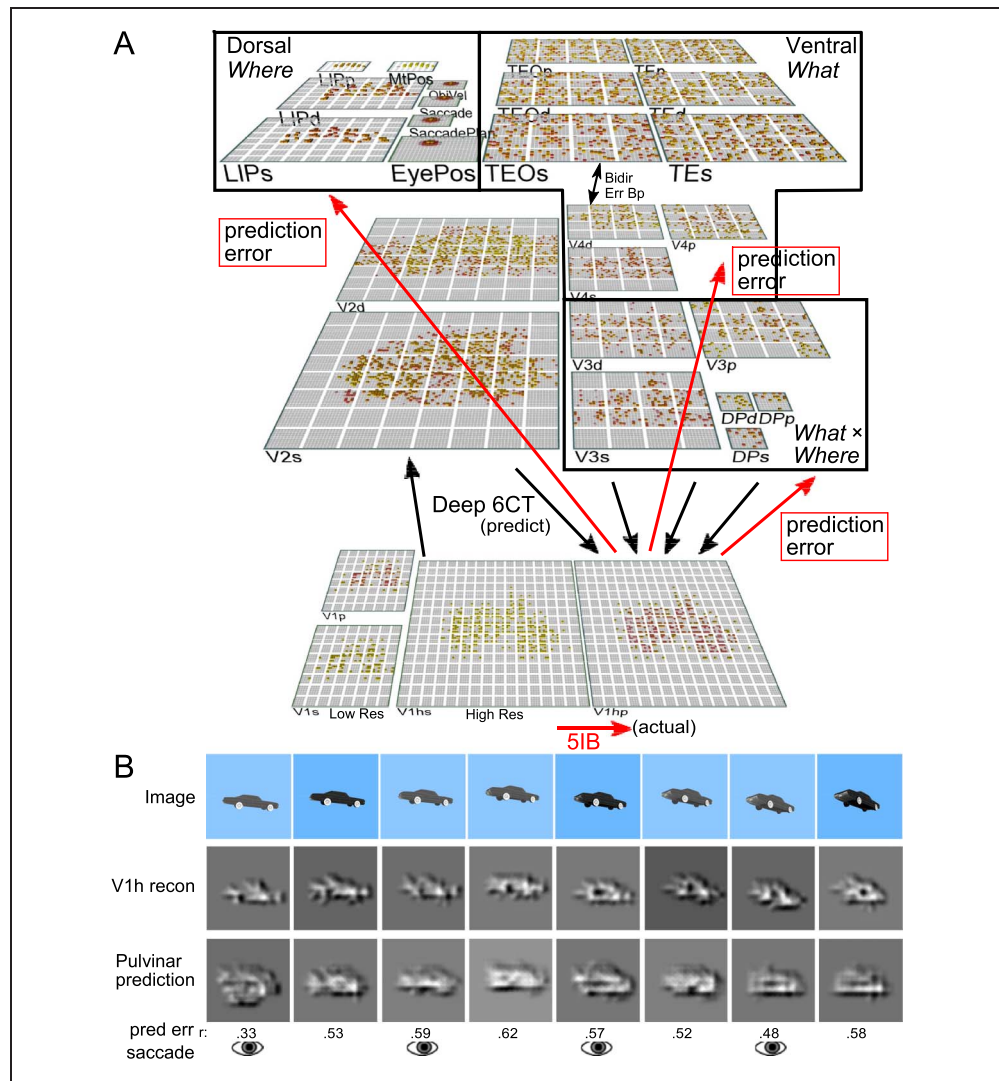


Figure 4. Predictive learning model applied to the FSA grammar shown in the previous figure. The first three panels (A–C) show the prediction state (end of the minus phase, e.g., the first 75 msec of an alpha cycle) of the trained model on the first three steps of the sequence “BTX” (plus phases also occurred but are not shown). The last panel (D) shows the plus phase after the third step. The “Input” layer provides the 5IB drivers for the corresponding HiddenP pulvinar layer, so the plus phase is always based on the specific randomly selected path taken. The “Targets” layer is purely for display, showing the two valid possible labels that could have been predicted. To track learning, the model’s prediction is scored as accurate if either or both targets are activated. Computationally, the model is similar to an SRN, where the CT layer that drives the prediction over the pulvinar encodes the activation state from the previous time step (alpha cycle), because of the phasic bursting of the 5IB neurons that drive CT updating. Note how the CT layer in B reflects the “Hidden” activation state in A and likewise for C reflecting B. This is evident because we are using one-to-one connectivity between Hidden and HiddenCT layers (which works well in general, along with full lateral connectivity within the CT layer). Thus, although the correct answer is always present on the Input layer for each step, the CT layer is nevertheless attempting to predict this input based on the information from the prior time step. (A) In the first step, the “B” label is unambiguous and easily predicted (based on prior “E” context). (B) In the second step, the network correctly guesses that the “T” label will come next, but there is a faint activation of the other “P” alternative, which is also activated sometimes based on prior learning history and associated minor weight tweaks. (C) In the third step, both “S” and “X” are equally predicted. (D) In the plus phase, only the Input pattern (“X” on this trial) drives HiddenP activations, and the projections from the pulvinar back to the cortex convey both the minus-phase prediction and the plus-phase actual input. You can see one HiddenCT neuron, just above the arrow, visibly changes its activation as a result (and all neurons experience smaller changes), and learning in all these cortical (Hidden) layer neurons is a function of their local temporal difference between the minus and plus phases.

Figure 5. (A) The WWI deep predictive learning model. The dorsal “Where” pathway learns first, using easily abstracted spatial blobs, to predict object location based on prior motion, visual motion, and saccade efferent copy signals. This drives strong top-down inputs to lower areas with accurate spatial predictions, leaving the residual error concentrated on “What” and “What × Where” integration. The V3 and DP (dorsal prelunate) constitute the What × Where integration pathway, binding features and locations. V4, TEO, and TE are the What pathway, learning abstracted object category representations, which also drive strong top-down inputs to lower areas. Suffixes: s = superficial; d = deep; and p = pulvinar. (C) Example sequence of eight alpha cycles that the model learned to predict, with the reconstruction of each image based on the V1 Gabor filters (“V1h recon”) and model-generated prediction (correlation r prediction error shown). The low resolution (Res) and reconstruction distortion impair visual assessment, but r values are well above the r s for each V1 state compared to the previous time step (mean = .38, minimum of .16 on Frame 4; see Appendix for more analyses). Eye icons indicate when a saccade occurred.



predicted to see in the V1 input in the next frame and what was actually seen.

As described in detail in the Appendix, our model was constructed to capture critical features of the visual system, including the major division between a dorsal “where” pathway and a ventral “what” pathway (Ungerleider & Mishkin, 1982), and the overall hierarchical organization of these pathways derived from detailed connectivity analyses (Markov, Ercsey-Ravasz, et al., 2014; Markov, Vezoli, et al., 2014; Felleman & Van Essen, 1991; Rockland & Pandya, 1979). In addition to these biological constraints, we conducted extensive exploration of the connectivity and architecture space and found a remarkable convergence between what worked functionally and the known properties of these pathways (O’Reilly et al., 2017). For example, the feedforward pathway has projections from lower-level superficial layers to superficial layers of higher levels, whereas feedback originated in both the superficial and deep layers and projected back to both (Felleman & Van Essen, 1991; Rockland & Pandya, 1979). In addition,

consistent with the core features of the pulvinar pathways discussed above, deep layer predictive (6CT) inputs originated in higher levels, whereas driver (5IB) inputs originated in lower levels. For simplicity, we organized the model layers in terms of these driver inputs, whereas the topographic organization of pulvinar in the brain is organized more according to the 6CT projection loops (Shipp, 2003).

Another important set of parameters are the strength of deep-layer recurrent projections, which influence the time-scale of temporal integration, producing a simple biologically based version of slow feature analysis (Wiskott & Sejnowski, 2002; Foldiak, 1991). We followed the biological data suggesting that recurrence increases progressively up the visual hierarchy (Chaudhuri, Knoblauch, Gariel, Kennedy, & Wang, 2015). It was essential that the “where” pathway learn first, consistent with extant data (Kiorpes, Price, Hall-Haro, & Movshon, 2012; Bourne & Rosa, 2006), including early pathways interconnecting lateral inferior parietal (LIP) and pulvinar (Bridge, Leopold, & Bourne, 2016), and a rare asymmetric pathway, from V1 to LIP (Markov, Ercsey-Ravasz, et al.,

2014), providing a direct shortcut for high-level spatial representations in LIP. Results from various informative model architecture and parameter manipulations are discussed below after the primary results from the standard intact model.

Learning curves and other model details are shown in the Appendix. We have also implemented a full de-novo replication of the model in a new modeling framework, which also replicated the results shown here. Furthermore, much of the model was originally developed in the context of a set of object-like patterns generated systematically from a set of simple line features (O'Reilly et al., 2017), and the parameters that work best in terms of combinatorial generalization on those patterns also worked well for these 3-D objects. Thus, we are confident that the model's learning behavior is not idiosyncratic to the particular set of objects used here and represents a general capacity of the system to develop abstract representations through predictive learning. Other ongoing work to be reported in an upcoming publication is applying the model to prediction of auditory speech inputs, which has a natural temporal structure, and finding similar results in terms of learning higher-level abstract encoding of these auditory signals.

To directly address the question of whether the hierarchical structure of the network supports the development of abstract, higher-level representations that go beyond the information present in the visual inputs, we applied a second-order similarity measure across the object-level similarity matrices computed at each layer in the network (Figure 6). This shows the extent to which the similarity matrix across objects in one layer is itself similar to the object similarity matrix in another layer, in terms of a correlation measure across these similarity matrices. Critically, this measure does not depend on any kind of subjective




Figure 6. Emergence of abstract category structure over the hierarchy of layers, comparing similarity structure in each layer versus that present in V1 (black line) or in TE (red line). Both cases, which are roughly symmetric, clearly show that IT layers (TEO, TE) progressively differentiate from raw input similarity structure present in V1 and, critically, that the model has learned structure beyond that present in the input. This is the simplest, most objective summary statistic showing this progressive emergence of structure, whereas subsequent figures provide a more concrete sense of what kinds of representations actually developed. *correl* = correlation.

interpretation of the learned representations—it only tells us whether whatever similarity structure was learned differs across the layers. Starting from either V1 compared to all higher layers, or the highest TE layer compared to all lower layers, we found a consistent pattern of progressive emergence of the object categorization structure in the upper IT pathway (TEO, TE).

This analysis confirms that indeed the IT category structure is significantly different from that present at the level of the V1 primary visual input. Thus, the model, despite being trained only to generate accurate visual input-level predictions, has learned to represent these objects in an abstract way that goes beyond the raw input-level information. We further verified that, at the highest IT levels in the model, a consistent, spatially invariant representation is present across different views of the same object (e.g., the average correlation across frames within an object was .901).

To better understand the nature of these learned representations, Figure 7 shows a representational similarity analysis (RSA) on the activity patterns at the highest IT layer (TE), which reveals the explicit categorical structure of the learned representations (Cadieu et al., 2014; Kriegeskorte, Mur, & Bandettini, 2008). Specifically, we found that the highest IT layer (TE) produced a systematic organization of the 156 3-D objects into five categories. In our admittedly subjective judgment, these categories seemed to correspond to the overall shape of the objects, as shown by the object exemplars in Figure 7 (pyramid shaped, vertically elongated, round, boxy/square, and horizontally elongated). Furthermore, the basic-level categories were subsumed within these broader shape-level categories, so the model appears to be sensitive to the coherence of these basic-level categories as well, but apparently, their shapes were not sufficiently distinct between categories to drive differentiated TE-level representations for each such basic-level category.

Given that the model only learns from a passive visual experience of the objects, it has no access to any of the richer interactive multimodal information that people and animals would have. Furthermore, as evident in Figure 5B, the relatively low resolution of the V1 layers (required to make the model tractable computationally) means that complex visual details are not reliably encoded (and, even so, are not generally reliable across object exemplars), such that the overall object shape is the most salient and sensible basis for categorization for this model.

Although these object shape categories appeared sensible to us, we ran a simple experiment to test whether a sample of 30 human participants would use the same category structure in evaluating the pairwise similarity of these objects. Figure 7B shows the results, confirming that indeed this same organization of the objects emerged in their similarity judgments. These judgments were based on the V1 reconstruction as shown in Figure 5B to capture the model's coarse-grained perception (see Appendix for methods and further analysis).

The progressive emergence of increasingly abstract category structure across visual areas, evident in Figure 6, has

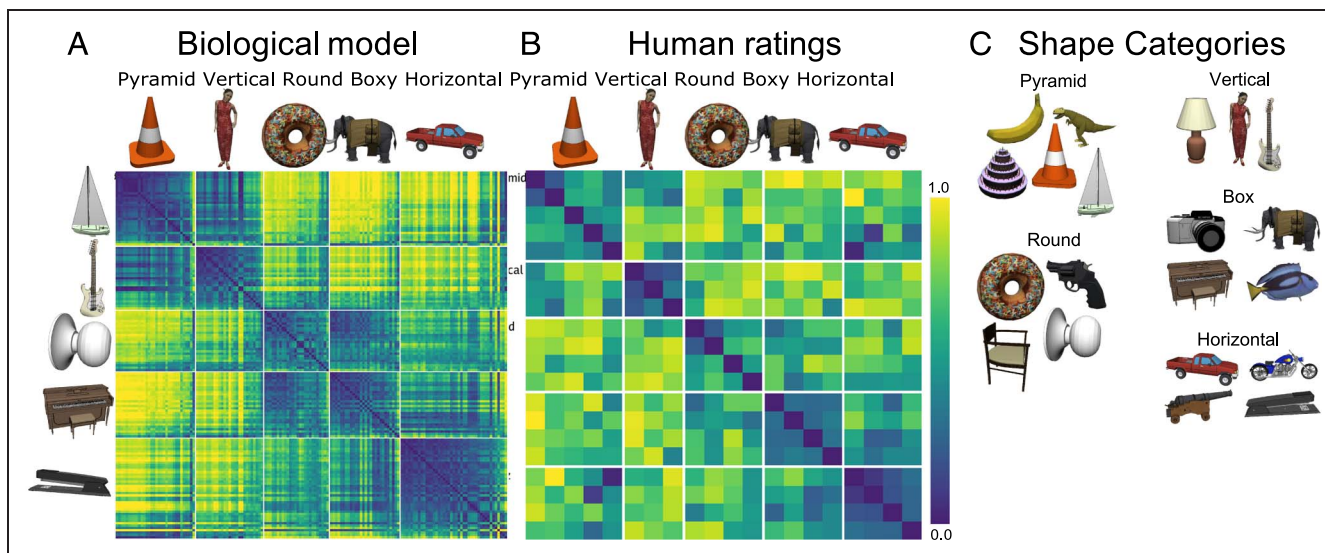


Figure 7. (A) Category similarity structure that emerged in the highest layer, TE, of the biologically based predictive learning model, showing dissimilarity (1-correlation) of the TE representation for each 3-D object against every other 3-D object (156 total objects). Blue cells have high similarity. Model has learned block-diagonal clusters or categories of high-similarity groupings, contrasted against dissimilar off-diagonal other categories. Clustering maximized average within-between dissimilarity (see Appendix) and clearly corresponded to the shape-based categories, with exemplars from each category shown. In addition, all items from the same basic-level object categories ($n = 20$) are reliably subsumed within learned categories. (B) Human similarity ratings for the same 3-D objects, presented with the V1 reconstruction (see Figure 1B) to capture coarse perception in the model, aggregated by 20 basic-level categories (a 156×156 matrix was too large to sample densely experimentally). Each cell is $1 -$ proportion of time given that object pair was rated more similar than another pair (see Appendix). The human matrix shares the same centroid categorical structure as the model (confirmed by permutation testing and agglomerative cluster analysis; see Appendix), indicating that human raters used the same shape-based category structure. (C) One object from each of the 20 basic-level categories, organized into the shape-based categories. The Vertical, Box, and Horizontal categories are fairly self-evident, and the model was most consistent in distinguishing those, along with subsets of the Pyramid (layered cake, traffic cone, sailboat) and Round (donut, doorknob) categories, whereas banana, trex, chair, and handgun were more variable.

been investigated in recent comparisons between monkey electrophysiological recordings and DCNNs, which provide a reasonably good fit of the overall progressive pattern of increasingly categorical organization (Cadieu et al., 2014). However, these DCNNs were trained on large data sets of human-labeled object categories, and it is perhaps not too surprising that the higher layers closer to these category output labels exhibited a greater degree of categorical organization. In contrast, because the only source of learning in our model comes from prediction errors over the V1 input layers, the graded emergence of an object hierarchy here reflects a truly self-organizing learning process.

Figure 8 compares the similarity structures in Layers V4 and IT in macaque monkeys (Cadieu et al., 2014) with those in corresponding layers in our model. In both the monkeys and our model, the higher IT layer builds upon and clarifies the noisier structure that is emerging in the earlier V4 layer, showing that our model replicates the essential qualitative hierarchical progression in the brain. As noted, we would not expect our model to exactly replicate the detailed object-specific similarity structure found in macaques, because of the impoverished nature of our model's experience, so this comparison remains qualitative in terms of the respective differences between V4 and IT in each model, rather than a direct comparison of the similarity structure between corresponding layers in the model and the macaque. In the future, when we can

scale up our model and tune the attentional processing dynamics necessary to deal with cluttered visual scenes, we will be able to train our model on the same images presented to the macaques and can provide this more direct comparison.

Finally, we did not use analyses based on decoding techniques, because with high-dimensional distributed neural representations, it is generally possible to decode many different features that are not otherwise compactly and directly represented (Fusi, Miller, & Rigotti, 2016). In preliminary work using decoding in the context of the simpler feature-based input patterns, we indeed found that decoding was not a very sensitive measure of the differentiation of representations across layers, which is so clearly evident in Figure 6. Thus, as advocates of the RSA approach have argued, measuring similarity structure evident in the activity patterns over a given layer generally provides a clearer picture of what that layer is explicitly encoding (Kriegeskorte et al., 2008).

In summary, the model learned an abstract category organization that reflects the overall visual shapes of the objects as judged by human participants, in a way that is invariant to the differences in motion, rotation, and scaling that are present in the V1 visual inputs. We are not aware of any other model that has accomplished this signature computation of the ventral "what" pathway in a purely self-organizing manner operating on realistic 3-D visual

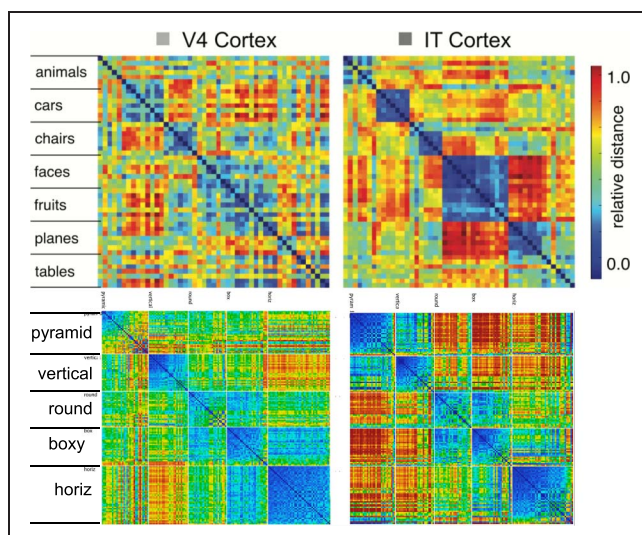
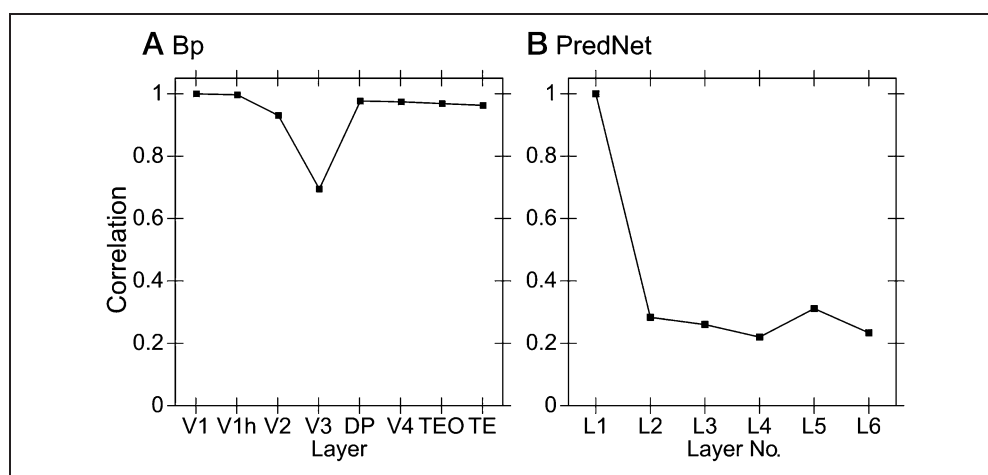


Figure 8. Comparison of progression from V4 to IT in macaque monkey visual cortex (top row, from Cadieu et al., 2014) versus same progression in model (replotted using a comparable color scale). Although the underlying categories are different, and the monkeys have a much richer multimodal experience of the world to reinforce categories such as foods and faces, the model nevertheless shows a similar qualitative progression of stronger categorical structure in IT, where the block-diagonal highly similar representations are more consistent across categories and the off-diagonal differences are stronger and more consistent as well (i.e., categories are also more clearly differentiated). Note that the critical difference in our model versus those compared in Cadieu et al. (2014) and related papers is that they explicitly trained their models on category labels, whereas our model is entirely self-organizing and has no external categorical training signal. horiz = horizontal.

objects, without any explicit supervised category labels. Furthermore, our model does this using a learning algorithm directly based on detailed properties of the underlying biological circuits in this pathway, providing a coherent overall account.

Figure 9. Similarity of similarity structure across layers for the comparison Bp models, comparing each layer to the first layer. (A) Bp model with the same “what/where” structure as the biological model. Unlike the biologically based model (Figure 6), the higher IT layers (TE, TEO) do not diverge significantly from the similarity structure present in V1, indicating that the model has not developed abstractions beyond the structure present in the visual input. Layer V3 is most directly influenced by spatial prediction errors, so it differs from both in strongly encoding position information. (B) PredNet model, which has six layers. Layers 2–6 diverge from Layer 1, but there is no progressive change in the higher layers as we see in our model moving from V4 to TEO. The divergence in correlation starting at Layer 2 is likely because of the fact that higher layers only encode errors, not stimulus-driven positive representations of the input. Aside from this large distinction (which is inconsistent with the similarity in neural coding seen in actual V1 and V2 recordings), there is no evidence of a cumulative development of abstraction in higher layers.



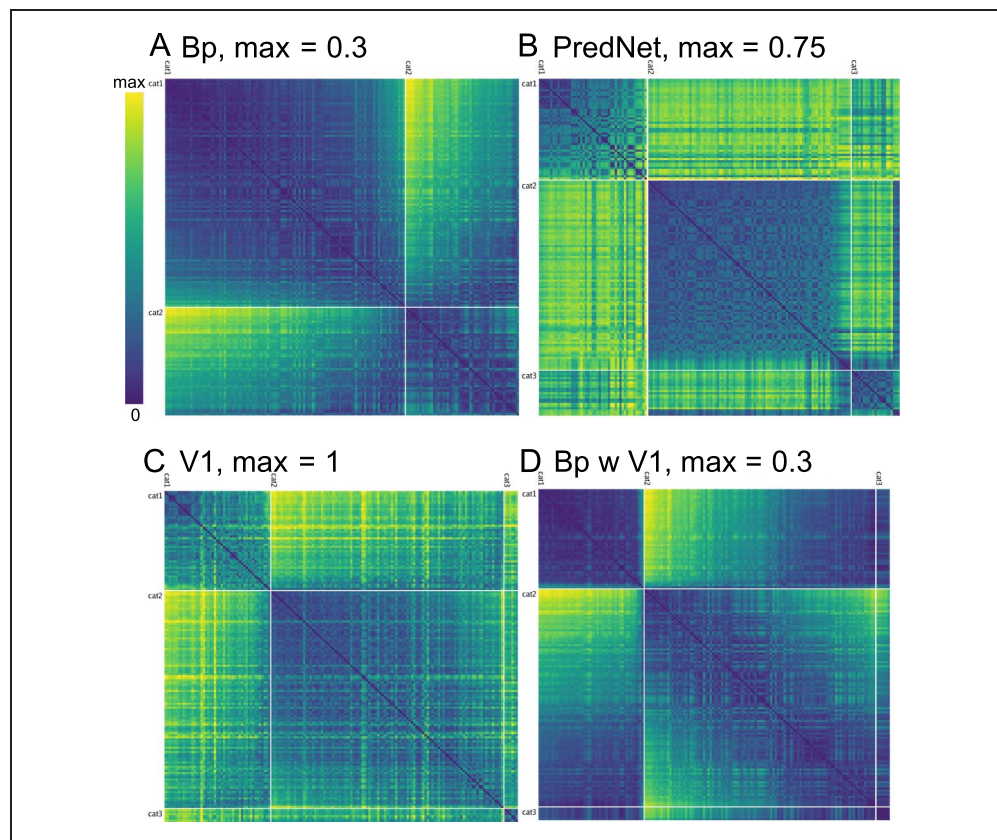
Backpropagation Comparison Models

To help discern some of the factors that contribute to the categorical learning in our model and provide a comparison with more widely used error Bp models, we tested a Bp-based version of the same “what vs. where” architecture as our biologically based predictive error model, and we also tested a standard PredNet model (Lotter et al., 2016) with extensive hyperparameter optimization (see Appendix). Because of the constraints of Bp, we had to eliminate any bidirectional connectivity loops in the Bp version, but we were able to retain a form of predictive learning by configuring the V1p pulvinar layer as the final target output layer, with the target being the next visual input relative to the current V1 inputs.

Figure 9 shows the same second-order similarity analysis as Figure 6, to determine the extent to which these comparison networks also developed more abstract representations in the higher layers that diverge from the similarity structure present in the lowest layers. According to this simple objective analysis, they did not—the higher layers showed no significant, progressive divergence in their similarity structure. The PredNet model did show a larger difference between the first layer and the rest of the layers, because of the subsequent layers encoding errors while the first layer has a positive representation of the image, but there was no progressive difference beyond that up into the higher layers.

Next, we examined the RSA matrices for the highest (TE) layer in the comparison models, also in comparison with the same for the V1 layer (Figure 10). This shows that the TE layer in the Bp model formed a simple binary category structure overall, which is similar to the RSA for the V1 input layer. It is also important to emphasize that the scales on these figures are different (as shown in their headers), such that these comparison models had much

Figure 10. (A) Best-fitting category similarity for TE layer of the Bp model with the same “what/where” structure as the biological model. Only two broad categories are evident, and the lower maximum (max) distance (0.3 vs. 1.5 in the biological model) means that the patterns are much less differentiated overall. (B) Best-fitting similarity structure for the PredNet model, in the highest of its layers (Layer 6), which is more differentiated than Bp (max = 0.75) but also less cleanly similar within categories (i.e., less solidly blue along the block diagonal) and overall follows a broad category structure similar to V1. (C) The best-fitting V1 structure, which has two broad categories and banana, is in a third category by itself. The lack of dark blue on the block diagonal indicates that these categories are relatively weak, and every item is fairly dissimilar from every other. (D) The Bp TE similarity values from A shown in the same ordering as V1 from C, demonstrating how the similarity structure has not diverged very much, consistent with the results shown in Figure 9—the within-between contrast differences are 0.0838 for A and 0.0513 for D (see Appendix for details).



less differentiated representations overall. Similar results were found in the PredNet model. Because existing work with these models has typically relied on additional supervised learning and decoder-based analyses (which are essentially equivalent to an additional layer of supervised learning), these RSA-based analyses provide an important, more sensitive way of determining what they learn purely through predictive learning.

These results show that the additional biologically derived properties in our model are playing a critical role in the development of abstract categorical representations that go beyond the raw visual inputs. These properties include excitatory bidirectional connections, inhibitory competition, and an additional Hebbian form of learning that serves as a regularizer (similar to weight decay) on top of predictive error-driven learning (O’Reilly & Munakata, 2000; O’Reilly, 1998). Each of these properties could promote the formation of categorical representations. Bidirectional connections enable top-down signals to consistently shape lower-level representations, creating significant attractor dynamics that cause the entire network to settle into discrete categorical attractor states. Another indication of the importance of bidirectional connections is that a greedy layer-wise pretraining scheme, consistent with a

putative developmental cascade of learning from the sensory periphery on up (Valpola, 2015; Bengio, Yao, Alain, & Vincent, 2013; Hinton & Salakhutdinov, 2006; Shrager & Johnson, 1996), did not work in our model. Instead, we found it essential that higher layers, with their ability to form more abstract, invariant representations, interact and shape learning in lower layers right from the beginning.

Furthermore, the recurrent connections within the TEO and TE layers likely play an important role by biasing the temporal dynamics toward longer persistence (Chaudhuri et al., 2015). By contrast, Bp networks typically lack these kinds of attractor dynamics, and this could contribute significantly to their relative lack of categorical learning. Hebbian learning drives the formation of representations that encode the principal components of activity correlations over time, which can help more categorical representations coalesce (and results below already indicate its importance). Inhibition, especially in combination with Hebbian learning, drives representations to specialize on more specific subsets of the space.

Ongoing work is attempting to determine which of these is essential in this case (perhaps all of them) by systematically introducing some of these properties into the Bp model, although this is difficult because full bidirectional recurrent

activity propagation, which is essential for conveying error signals top-down in the biological network, is incompatible with the standard efficient form of error Bp, and requires significantly more computationally intensive and unstable forms of fully recurrent Bp (Williams & Zipser, 1992; Pineda, 1987). Furthermore, Hebbian learning requires dynamic inhibitory competition, which is difficult to incorporate within the Bp framework.

Architecture and Parameter Manipulations

Figure 11 shows only a few of the large number of parameter manipulations that have been conducted to develop and test the final architecture. For example, we hypothesized that separating the overall prediction problem between a spatial “where” versus nonspatial “what” pathway (Goodale & Milner, 1992; Ungerleider & Mishkin, 1982) would strongly benefit the formation of more abstract, categorical object representations in the “what” pathway. Specifically, the “where” pathway can learn relatively quickly to predict the overall spatial trajectory of the object (and anticipate the effects of saccades) and thus effectively regress out that component of the overall prediction error, leaving the residual error concentrated in object feature information, which can train the ventral “what” pathway to develop abstract visual categories.

Figure 11A shows that, indeed, when the “where” pathway is lesioned, the formation of abstract categorical representations in the intact “what” pathway is significantly impaired. We also hypothesized that full predictive learning (about the future), as compared to just encoding and decoding the current state (i.e., an autoencoder, which is much easier computationally), is also critical for the formation of abstract categorical representations—prediction is a “desirable difficulty” (Bjork, 1994). Figure 11B shows that this was the case. Finally, consistent with our

hypothesis that Hebbian learning provides an important bias on learning, Figure 11C shows the impairment associated with reducing this learning bias. The significant reduction in differentiation across all of these manipulations shows that this differentiation property is not a simple consequence of the neural architecture but rather depends critically on the learning process, unfolding over time with appropriate parameter values and other architectural components. Furthermore, the Bp comparison model shares the same architecture and does not show the differentiation across layers.

Predictive Behavior

A signature example of predictive behavior at the neural level in the brain is the predictive remapping of visual space in anticipation of a saccadic eye movements (Marino & Mazer, 2016; Nakamura & Colby, 2002; Gottlieb, Kusunoki, & Goldberg, 1998; Colby, Duhamel, & Goldberg, 1997; Duhamel et al., 1992; Figure 12A). Here, parietal neurons start to fire at the future receptive field location where a currently visible stimulus will appear after a planned saccade is actually executed. Remapping has also been shown for border ownership neurons in V2 (O’Herron & von der Heydt, 2013) and in Area V4 (Neupane, Guitton, and Pack, 2016, 2020). These are examples, we believe, of a predictive process operating throughout the neocortex to predict what will be experienced next. A major consequence of this predictive process is the perception of a stable, coherent visual world despite constant saccades and other sources of visual change.

Figure 12B shows that our model exhibits this predictive remapping phenomenon. Specifically, LIP, which is most directly interconnected with the saccade efferent copy signals, is the first to predict the new location, and it then drives top-down activation of lower layers. This top-down

Figure 11. Effects of various manipulations on the extent to which TE representations differentiate from V1. For all plots, “Intact” is the same result shown in Figure 6 from the intact model for ease of comparison (A is missing V3 and DP dorsal pathway layers). All of the following manipulations significantly impair the development of abstract TE categorical representations (i.e., TE is more similar to V1 and the other layers). (A) Dorsal “where” pathway lesions, including lateral inferior parietal sulcus (LIP), V3, and dorsal prelunate (DP). This pathway is essential for regressing out location-based prediction errors, so that the residual errors concentrate feature-encoding errors that train the “what” pathway. (B) Allowing the deep layers full access to current-time information, thus effectively eliminating the prediction demand and turning the network into an autoencoder, which significantly impairs representation development and supports the importance of the challenge of predictive learning for developing deeper, more abstract representations. (C) Reducing the strength of Hebbian learning by 20% (from 2.5 to 2), demonstrating the essential role played by this form of learning on shaping categorical representations. Eliminating Hebbian learning entirely (not shown) prevented the model from learning anything at all, as it also plays a critical regularization and shaping role on learning. Hebb = Hebbian.

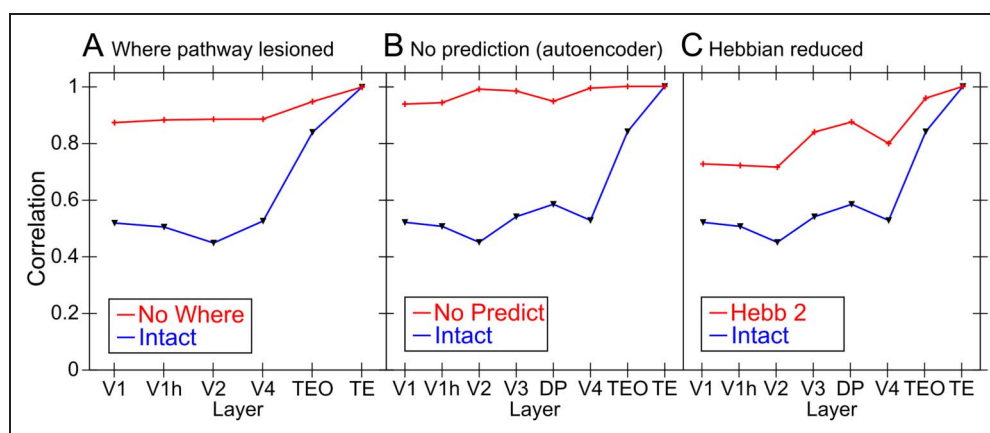
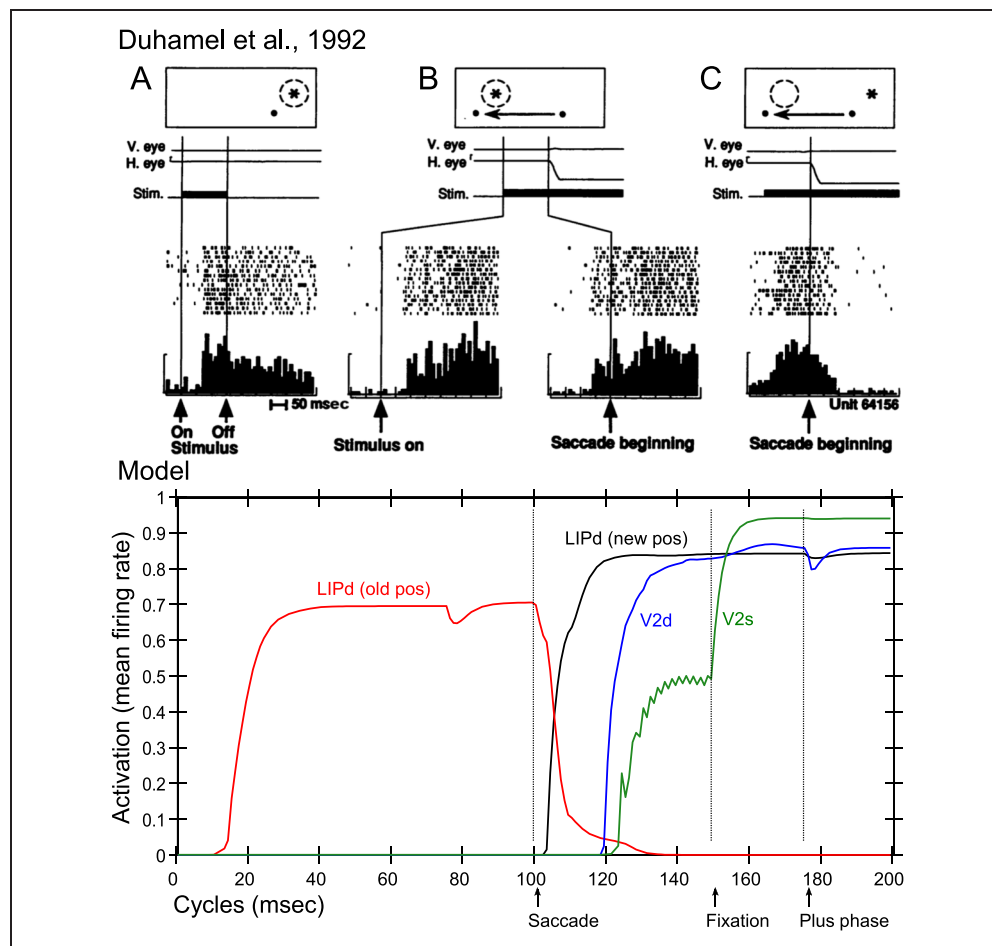


Figure 12. Predictive remapping. (Top) Original remapping data in LIP from Duhamel et al. (1992). A shows stimulus (star) response within the receptive field (dashed circle) relative to the fixation dot (top right of fixation). (B) Just before monkey making a saccade to new fixation (moving left), stimulus is turned on in receptive field location that will be upper right of the new fixation point, and the LIP neuron responds to that stimulus in advance of the saccade completing. The neuron does not respond to the stimulus in that location if it is not about to make a saccade that puts it within its receptive field (not shown). This is predictive remapping. (C) Response to the old stimulus location goes away as saccade is initiated. (Bottom) Data from our model, from individual units in LIPd, V2d, and V2s, showing that the LIP deep neurons respond to the saccade first, activating in the new location and deactivating in the old, and this LIP activation goes top-down to V3 and V2 to drive updating there, generally at a longer latency and with less activation especially in the superficial layers. When the new stimulus appears at the point of fixation (after a 50-msec saccade here), the primed V2s units get fully activated by the incoming stimulus. However, the deep neurons are insulated from this superficial input until the plus phase, when the cascade of 5IB firing drives activation of the actual stimulus location into the pulvinar, which then reflects up into all the other layers.



dynamic is consistent with the account of predictive remapping given by Wurtz (2008) and Cavanagh et al. (2010), who argue that the key remapping takes place at the high levels of the dorsal stream, which then drive top-down activation of the predicted location in lower areas, instead of the alternative where lower levels remap themselves based on saccade-related signals. The lower-level visual layers are simply too large and distributed to be able to remap across the relevant degrees of visual angle—the extensive lateral connectivity needed to communicate across these areas would be prohibitive.

NEURAL DATA AND PREDICTIONS

Having tested the computational and functional learning properties of this biologically based predictive learning mechanism, we now return to consider some of the most important neural data of relevance to our hypotheses, beyond that summarized in the introduction, including contrasts with a widely discussed alternative framework for predictive coding, and some of the extensive data on

alpha frequency effects, followed by a discussion of predictions that would clearly test the validity of this framework.

Additional Neuroscience Data

We begin with data relevant to the basic neural-level properties of the framework. First, a central element of the proposed model is the alpha cycle bursting, and subsequent interburst pauses, in the 5IB neurons. Direct electrophysiological recording of deep layer neurons shows periodic alpha-scale bursting for continuous tones in awake animals (Luczak et al., 2009, 2013; Sakata & Harris, 2009, 2012). In vitro, a variety of potential mechanisms behind the generation and synchronization of the 5IB bursts driving this alpha cycle have been identified (Franceschetti et al., 1995; Silva et al., 1991; Connors et al., 1982). Furthermore, the pulvinar has been shown to drive alpha-frequency synchronization of cortical activity across areas in the alpha band in awake, behaving animals (Saalmann et al., 2012). We review the larger alpha frequency literature in more detail below, but it is critical to emphasize that this alpha bursting dynamic is

actually found in awake, behaving animals, because so many other bursting and up/down state phenomena have recently been shown to only occur in anesthetized brains, including bursting in the thalamic TRC neurons.

In contrast to the 5IB bursting, the 6CT neurons exhibit regular spiking behavior (Thomson, 2010; Thomson & Lamy, 2007), providing consistent activation to the pulvinar. In addition, they do not have axonal branches that project to other cortical areas—the subpopulation that projects to the pulvinar only project there and not to other cortical areas (Petrof, Viaene, & Sherman, 2012), whereas there are other Layer 6 neurons that do project to other cortical areas. This distinct connectivity is consistent with a specific role of this neuron type in generating predictions in the pulvinar. The 6CT synaptic inputs on pulvinar TRCs have metabotropic glutamate receptors that have longer timescale temporal dynamics consistent with the alpha period (100 msec) and even longer (Sherman, 2014), and the 6CT neurons themselves also have temporally delayed responding (Harris & Shepherd, 2015; Thomson, 2010; Sakata & Harris, 2009). Furthermore, they have significantly more plasticity-inducing *N*-methyl-D-aspartate receptors compared to the 5IB projections (Usrey & Sherman, 2018). These properties are consistent with the 6CT inputs driving a longer-integrated prediction signal that is subject to learning, whereas the 5IB are likely nonplastic, and their effects are tightly localized in time.

The 5IB inputs often have distinctive glomeruli structures at their synapses onto pulvinar neurons, which contain a

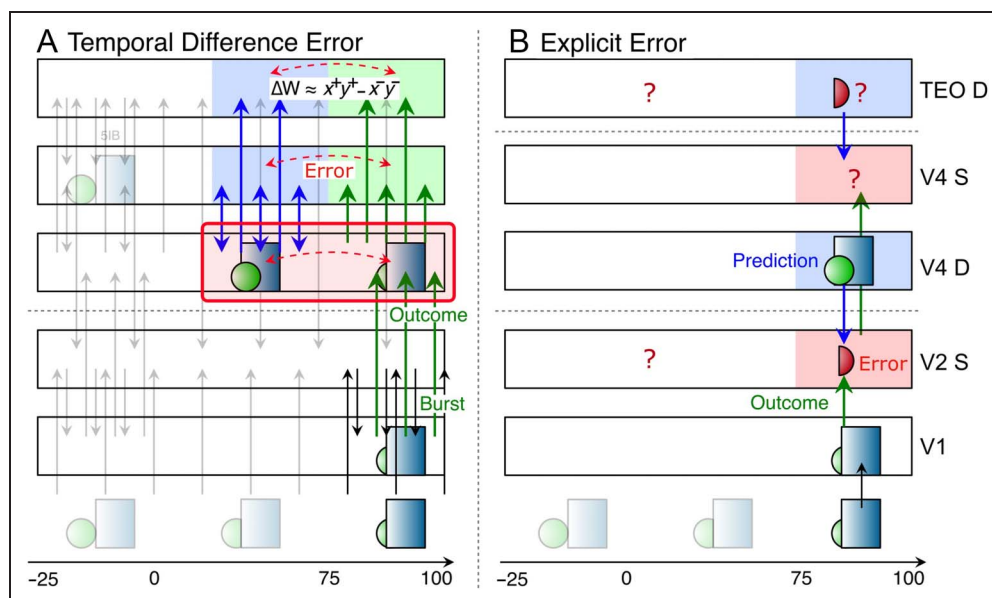
complete feedforward inhibition circuit involving a local inhibitory interneuron, in addition to the direct strong excitatory driver input (Wilson, Bose, Sherman, & Guillery, 1984). Computationally, this can provide a balanced level of excitatory and inhibitory drive so as to not overly excite the receiving neuron, while still dominating its firing behavior.

Although there are well-documented and widely discussed burst versus tonic firing modes in pulvinar neurons (Sherman & Guillery, 2006), there is not much evidence of these playing a clear role in the awake, behaving state, and as noted earlier, the growing electrophysiological evidence shows a remarkable correspondence between cortical and pulvinar response properties across multiple different pulvinar areas in this awake state. Nevertheless, there may be important dynamics arising from these firing modes that are more subtle or emerge in particular types of state transitions that may have yet to be identified.

Contrast with Explicit Error Frameworks

To further clarify the nature of the present theory and introduce a body of relevant data, we contrast it with the widely discussed EE framework for predictive coding (Lotter et al., 2016; Bastos et al., 2012; Oudén et al., 2012; Friston, 2005, 2010; Rao & Ballard, 1999; Kawato et al., 1993; Figure 13). The hypothesized locus for computing errors in this framework is in the superficial layers of the neocortex, which are suggested to directly compute the difference between

Figure 13. Comparison between (A) the proposed thalamocortical temporal-difference predictive learning model (from Figure 2) versus (B) the Bayesian-style EE coding model (Bastos et al., 2012; Friston, 2010; Rao & Ballard, 1999). The EE model holds that superficial (S, lamina 2/3) error-coding neurons receive the prediction via a net inhibitory top-down projection from higher-level deep layer (D) neurons and an excitatory bottom-up projection representing the outcome, such that their activation represents the difference. To encode both signs of the error (omissions, false alarms) with positive-only spike rates, two separate populations of EE neurons would be required, or a more complicated deviation from tonic firing level scheme.



Unambiguous evidence of such EE coding neurons has not been found (Walsh et al., 2020). In contrast, error signals in our proposed framework remain as a temporal difference between the two states of prediction versus outcome, which enables all connectivity between cortical areas to be excitatory and always represent a positive encoding of either the prediction or outcome. In contrast, under EE, after one error subtraction at the lowest level, only error signals are hypothesized to flow forward to higher layers, meaning that the representations at higher layers are about increasingly higher-order errors, not positive encodings of the environmental state at increasing levels of abstraction. These are indicated by ? because they are difficult to picture intuitively, and they are inconsistent with extensive available data showing similar positive representations of the external world at all levels in the visual hierarchy. Although some frameworks make claims about temporal dynamics, these are not strongly constrained by the basic computational framework, so that also remains a question.

bottom-up inputs from lower layers and top-down inputs from higher areas. Despite many attempts to identify such EE coding neurons in the cortex, no substantial body of unambiguous evidence has been discovered (Walsh, McGovern, Clark, & O'Connell, 2020; Kok & de Lange, 2015; Kok, Jehee, & de Lange, 2012; Summerfield & Egner, 2009; Lee & Mumford, 2003). Furthermore, because of the positive-only firing rate nature of neural coding, two separate populations would be required to convey both signs of prediction error signals, or it would have to be encoded as a variation from tonic firing levels, which are generally low in the neocortex.

By contrast, the use of temporal difference error signals enables all connections between cortical layers to be excitatory, and each layer can represent the positive encoding of either the prediction or outcome state, at different levels of abstraction. These properties are overwhelmingly supported by extensive electrophysiological data about the hierarchical organization of representations, for example, in the visual object recognition pathway (Cadieu et al., 2014; VanRullen & Thorpe, 2002; Kobatake & Tanaka, 1994), and are consistent with the widely supported biased competition model for excitatory top-down attentional effects (O'Reilly et al., 2013; Miller & Cohen, 2001; Reynolds et al., 1999; Desimone & Duncan, 1995).

The EE approach requires net inhibitory top-down predictions, and it sends error signals forward, not positive representations of the actual state at a given level of abstraction. Thus, a literal interpretation (and at least one existing implementation; Lotter et al., 2016) has only error signals represented at all levels above the lowest level, which is inconsistent with the positive encoding of stimuli at various levels of abstraction across the visual hierarchy. For example, although Issa, Cadieu, and DiCarlo (2018) observed an error-signal-like increase in activation for atypical faces in some posterior IT neurons, these neurons overall had a positive stimulus encoding, with only a relatively small, later, error-like modulation.

Furthermore, as discussed below, anticipatory predictions typically closely resemble the subsequent stimulus-driven activity, suggesting a positive, not inhibitory, effect (Walsh et al., 2020; Cavanagh et al., 2010; Lee & Mumford, 2003; Duhamel et al., 1992). However, there are various different ways of reformulating the neural implementation of EE that can avoid some of these issues (Bastos et al., 2012; Spratling, 2008), but perhaps, this flexibility renders the framework difficult to falsify (Kogo & Trengove, 2015). In any case, an extensive treatment of the issues with EE is beyond the scope of this paper and has already been aptly covered by Walsh et al. (2020)—our goal here is to highlight some of the core differences as a way to clarify the framework by way of contrast and in relation to available data.

First, there are many examples of anticipatory predictive neural firing in the brain. Of perhaps greatest relevance, Barczak et al. (2018) recently showed that the auditory

pulvinar in monkeys exhibits predictive firing using a carefully controlled auditory sequence that had no first-order acoustic differences from a background noise signal. The pulvinar predictive activation preceded that of A1, suggesting a strong predictive role for pulvinar. Unfortunately, the deep layers of higher auditory areas that should contribute to the formation of the pulvinar prediction were not recorded in this study, so their role in generating the prediction could not be determined.

Nevertheless, there is extensive additional evidence for top-down anticipatory activation of predicted stimuli, with activity patterns closely resembling the subsequent stimulus-driven ones (Walsh et al., 2020). For example, the widely replicated predictive remapping effect, simulated in our model (Figure 12), is of this nature (Cavanagh et al., 2010; Wurtz, 2008; Duhamel et al., 1992). The fact that these anticipatory activations are of a positive nature, consistent with the stimulus-driven activations, is inconsistent with the expected behavior of EE neurons, which should be inhibited by the top-down prediction, while not receiving any bottom-up stimulus.

However, the neural response to the actual predicted stimulus itself is typically suppressed relative to unexpected stimuli, that is, expectation suppression (Bastos et al., 2012; Meyer & Olson, 2011; Todorovic et al., 2011; Summerfield et al., 2008). This phenomenon is widely cited as evidence in favor of the EE predictive coding framework, consistent with an inhibitory effect of the expectation. Nevertheless, despite various conflicting results and many complications of interpretation, multiple comprehensive reviews conclude that it is difficult to distinguish expectation suppression from the neural adaptation effects that underlie the well-documented repetition suppression effect (Walsh et al., 2020; Kok & de Lange, 2015; Vinken & Vogels, 2017; Kok et al., 2012; Summerfield & Egner, 2009; Lee & Mumford, 2003). Furthermore, detailed single-neuron-level recordings are the least likely to show these effects—instead, they are most evident in aggregate signals such as the BOLD response in fMRI, suggesting that they may more strongly reflect population-level differences in activity, rather than individual EE coding neurons.

As noted earlier, accurately predicted outcomes in our framework would result in a continued adaptation of the neural response carrying over from the prediction to the outcome state, whereas unexpected outcomes would be associated with two distinct patterns of activity over a given area: first the prediction and then the outcome. Thus, the unexpected outcome state would not be subject to the prior neural adaptation effects, and furthermore, the time-integrated aggregate activity over these two patterns would be greater compared to the single activity state associated with an accurately predicted outcome. Thus, our model explains expectation suppression without invoking EE neurons, meaning that considerably more detailed and replicable experimental paradigms using single-neuron resolution techniques are needed to distinguish EE from our framework.

Alpha Frequency Effects

The alpha frequency bursting of 5IB neurons acting as drivers into the pulvinar naturally entrains the predictive learning process in our model to this fundamental rhythm, which has long been recognized as an important signature of posterior cortical function (VanRullen & Koch, 2003; Varela, Toro, John, & Schwartz, 1981; Nunn & Osselton, 1974; Walter, 1953; Berger, 1929). A number of different functional associations with alpha have been established, and this literature is large and growing rapidly. Thus, we refer the reader to recent reviews (Foster & Awh, 2019; Clayton et al., 2018; VanRullen, 2016; Jensen, Bonnefond, Marshall, & Tiesinga, 2015) while highlighting the data most relevant to our specific framework here, organized according to a set of key points.

- Alpha is specifically associated with deep neocortical layers and the pulvinar as well as with feedback pathways in the cortex. This has been established using direct laminar-specific electrophysiological single-neuron and local field potential recordings (Luczak et al., 2013; Spaak, Bonnefond, Maier, Leopold, & Jensen, 2012; Xing, Yeh, Burns, & Shapley, 2012; Buffalo et al., 2011; Maier, Aura, & Leopold, 2011; Maier, Adams, Aura, & Leopold, 2010) and feedforward versus feedback manipulations (Michalareas et al., 2016; Bastos et al., 2015; Jensen et al., 2015; van Kerkoerle et al., 2014; von Stein, Chiang, & König, 2000). These data are consistent with the 5IB alpha bursting and the major role of cortical deep layers in driving top-down corticocortical projections (in addition to the 6CT pathway that is specific to the pulvinar). By contrast, these same studies show that superficial cortical layers are associated with gamma frequency (40-Hz) dynamics. However, the next point raises some important interpretational difficulties.
- Increases in cortical activity levels, for example, because of attention, produce a corresponding decrease in alpha power, whereas decreased activity increases alpha power (Foster & Awh, 2019; Jensen & Mazaheri, 2010; Fries, Womelsdorf, Oostenveld, & Desimone, 2008; Klimesch, Sauseng, & Hanslmayr, 2007; Kelly, Lalor, Reilly, & Foxe, 2006; Worden, Foxe, Wang, & Simpson, 2000). This pattern is not exactly what you might expect if alpha was a signature of predictive learning. Furthermore, given that these same pulvinar and thalamocortical pathways are also widely regarded as important for attention (Fiebelkorn & Kastner, 2019; Zhou et al., 2016; Saalman & Kastner, 2011; Snow et al., 2009; Bender & Youakim, 2001; LaBerge & Buchsbaum, 1990), this pattern presents a challenge for many theorists. However, it is possible to explain this pattern as arising directly from the desynchronizing effects of cortical activity on alpha power. Specifically, neural spiking is associated with broadband noise, because of the highly random, Poisson nature of spike firing, which can desynchronize the entrainment of lower-frequency oscillations including alpha (Solomon et al., 2017; Privman, Malach, & Yeshurun, 2013; Waldert, Lemon, & Kraskov, 2013; Ray & Maunsell, 2011). In other words, because cortical activity is inherently noisy, it tends to interfere with the coherent activity across populations of neurons needed to produce a strong alpha frequency power signal. This explanation is directly supported by studies manipulating and measuring cortical activity (Zhou et al., 2016; Fries et al., 2008) and is consistent with alpha power changes being a result of attentional modulation, but not their cause (Antonov, Chakravarthi, & Andersen, 2020). Thus, although attention and predictive learning can both affect overall activity levels in the cortex and thus drive changes in alpha power, alpha power itself is not a transparent measure of the underlying mechanisms supporting these functions, which may help to explain some contradictory patterns of results (Gundlach, Moratti, Forschack, & Müller, 2020; Foster & Awh, 2019; Keitel et al., 2019).
- Alpha phase effects provide a more direct measure of thalamocortical function than alpha power and have been more consistently related to perception, attention, and prediction (Solís-Vivanco, Jensen, & Bonnefond, 2018; Neupane et al., 2017; Jaegle & Ro, 2013; Palva & Palva, 2011; Mathewson, Fabiani, Gratton, Beck, & Lleras, 2010; Busch, Dubois, & VanRullen, 2009; VanRullen & Koch, 2003; Varela et al., 1981; Nunn & Osselton, 1974). For example, weak, near-threshold stimuli are more reliably detected and processed when presented in the trough of the individual's ongoing alpha cycle. Of greatest relevance to this paper are studies showing effects of prediction on alpha phase (Mayer, Schwiedrzik, Wibral, Singer, & Melloni, 2016; Sherman, Kanai, Seth, & VanRullen, 2016; Samaha, Bauer, Cimaroli, & Postle, 2015). For example, Mayer et al. (2016) showed that prestimulus alpha phase directly correlated with the predictability of the upcoming stimulus, and the pattern of this prestimulus activation was indistinguishable from the subsequent stimulus activation pattern. This is consistent with our model, and less consistent with the EE framework, as discussed previously. Neupane et al. (2017) found strong alpha coherence effects in local field potential recordings distributed across V4, associated with the predictive remapping of receptive fields (Duhamel et al., 1992).
- Discrete, salient, or oscillatory stimuli entrain the alpha cycle in the brain (Spaak, de Lange, & Jensen, 2014; Mathewson et al., 2012). Furthermore, the massive literature on ERPs may represent a significant contribution from alpha-level entrainment (Klimesch, 2011; Gruber, Klimesch, Sauseng, & Doppelmayr, 2005; Makeig et al., 2002). These entrainment effects are consistent with the 5IB entrainment mechanisms in our framework, as described earlier, and entrainment is functionally important for aligning predictive learning with relevant salient or unexpected outcomes.
- The pulvinar contributes to synchronizing alpha phase relationships across different brain areas (Fiebelkorn,

Pinsk, & Kastner, 2018; Saalman et al., 2012). This is consistent with the broad, convergent pattern of projections into the pulvinar from many different cortical areas, and the corresponding broad projections back out to these same areas (Arcaro et al., 2015; Shipp, 2003). Functionally, this convergence and synchronization are important for integrating the contributions from these different areas at the same time, to generate predictions over the pulvinar.

- The theta cycle, composed of a pair of alpha cycles, organizes saccades as well as attentional, motor, and mnemonic processes (Fiebelkorn & Kastner, 2019). The theta rhythm is dominant in the medial temporal lobe and hippocampus and has been extensively studied there (Buzsáki, 2005; Kahana, Seelig, & Madsen, 2001). Furthermore, there is a sharp peak of saccade fixation durations at 200 msec, which suggests that two alpha cycles are typically required for complete processing of a given fixation. On the first cycle, the predictions from before the eye moved may be fairly vague depending on factors such as the size of the saccade and familiarity with the environment. However, after the first alpha cycle of a fixation, a subsequent postdiction phase can provide an important additional learning opportunity, to consolidate and more deeply encode the current fixation (computationally equivalent to an autoencoder). In addition, a mix of smaller saccades (including microsaccades) and larger saccades enables a range of more and less predictable outcomes on the first alpha cycle after the saccade and matches human behavior (Martinez-Conde, Otero-Millan, & Macknik, 2013; Martinez-Conde, Macknik, & Hubel, 2004).

Putting all of these points together, a particularly effective way of testing the predictions of our framework would be measuring alpha phase changes emerging in the pre-stimulus period as a function of predictive learning in predictable sequential stimulus streams. In addition, it would also be important to examine theta- and alpha-cycle dynamics in relation to predictive learning in the context of attention, motor control, and memory processes, to better understand the larger systems-level temporal organization of learning and processing in the brain (Fiebelkorn & Kastner, 2019).

Predictions for Predictive Learning

In this section, we enumerate a set of direct, testable predictions from our framework. Before doing so, there are several important considerations for any experimental test of the theory. First, the nature of what is to be learned must be matched to the pulvinar area in question. For example, learning a new variation of basic physics in movies at the alpha time scale (e.g., altering properties such as gravity, inertia, or elasticity) would be appropriate for the lower-level visual pathways. At higher visual levels

(e.g., IT cortex), it might be possible to use simple sequences of different objects, although it is not clear to what extent the hippocampus or PFC might also contribute in this case (Fiser et al., 2016; Gavornik & Bear, 2014). To distinguish pulvinar learning effects from pervasive motor learning supported by other brain areas, it would be most effective to directly measure activity in the pulvinar and/or associated perceptual neocortical areas, instead of involving overt behavioral performance.

Much of the learning in posterior sensory cortex should take place early in development, requiring very early developmental interventions or genetic knockouts that are expressed from the start (which can also have other interpretational issues if not highly selective). In our models, the bulk of the basic sensory predictive learning happens very quickly, because the basic first-level regularities are quite strong and relatively easily learned. Although there are longer-term changes in the higher-level pathways in our models, more fine-grained measurements would likely be required to see these changes. Once this learning has taken place, the remaining contributions of the thalamocortical circuit are likely more strongly weighted toward its role in attention, as we discuss below. Finally, directly lesioning or inactivating the pulvinar is not likely to be very informative, because existing work has shown dramatic effects on cortical activity (Zhou et al., 2016; Purushothaman, Marion, Li, & Casagrande, 2012), and furthermore, any effects could be attributed to the attentional contributions of the pulvinar.

With these considerations in mind, here are a set of strong predictions from our model that should be testable using existing techniques. Failure to obtain the predicted result, while adhering to all the relevant constraints, would constitute a falsification of our model.

- Blocking 5IB bursting mechanisms early in developmental learning should disrupt learning. It should be possible to selectively knock out or modify the channels that cause this specific population of neurons to burst fire, and doing so should have a significant effect on learning in associated neocortical and pulvinar areas, given the critical role that this burst firing plays on the predictive learning process, as elaborated above.
- Blocking synaptic plasticity in the pulvinar (specifically the 6CT inputs) very early in developmental learning should impair learning. Although most of the learning overall should occur in the neocortex as a result of the temporal difference error signal broadcast by the pulvinar (which should remain generally intact), learning in the 6CT projections is important, especially right at the start, to map the emerging neocortical representations into the space defined by the 5IB projections.
- Temporal differences on an alpha cycle timescale actually drive synaptic plasticity in an error-driven learning manner, in neocortical pyramidal neurons and in 6CT inputs to the pulvinar. That is, if a pre/post pair of neurons across a synapse is more active in the prediction

than the subsequent outcome, the synapse should experience long-term depression, and vice versa if the activity pattern is reversed (long-term potentiation, for more activity in outcome than prediction). Furthermore, if activity is essentially stable across both prediction and outcome phases, then weights should not change (modulo, a small level of Hebbian learning; O'Reilly et al., 2012; O'Reilly & Munakata, 2000). This should be directly testable using current experimental methods and is perhaps the single most important empirical test of this entire framework, and it also underlies many other current approaches to error-driven learning in the brain (Lillicrap et al., 2020; Whittington & Bogacz, 2019; Bengio et al., 2017). One general consideration is the extent to which an awake *in vivo* preparation would be required to capture all the neuromodulatory and other factors present when this learning normally takes place. Some suggestive evidence in such a preparation is generally consistent with a sensitivity to relatively short-term temporal dynamics (Lim et al., 2015), although these results lacked the direct measurement of individual neural activity across a synapse.

DISCUSSION

We have hypothesized a novel computational function for the distinctive features of thalamocortical circuits (Usrey & Sherman, 2018; Sherman & Guillery, 2006), as supporting a specific form of prediction-error driven learning, where predictions arise from the numerous top-down layer 6CT projections into the pulvinar, and the strong, sparse driving 5IB inputs supply the bottom-up sensory-driven outcome. The phasic bursting nature of the 5IB inputs results in a natural temporal-difference error signal of prediction followed by outcome, consistent with extensive neural recording data. This temporal dynamic is also essential for enabling predictions to be generated without contamination from current sensory inputs and predicts a characteristic alpha-frequency prediction cycle based on the 10-Hz bursting cycle of the 5IB inputs, consistent with the pervasive influence of alpha on perception and neural dynamics (Foster & Awh, 2019; Clayton et al., 2018; VanRullen, 2016; Jensen et al., 2015). In short, the hypothesized predictive learning function fits remarkably well with a number of well-established properties of these thalamocortical circuits, and we also provided a set of additional predictions that could be tested to further evaluate this theory, especially in contrast to the widely discussed alternative of EE coding neurons, which have not been unambiguously supported across a range of empirical studies (Walsh et al., 2020).

Furthermore, we implemented this theory in a large-scale model of the visual system and demonstrated that learning based strictly on predicting what will be seen next is, in conjunction with a number of critical biologically motivated network properties and mechanisms, capable

of generating abstract, invariant categorical representations of the overall shapes of objects. The nature of these shape representations closely matches human shape similarity judgments on the same objects. Thus, predictive learning has the potential to go beyond the surface structure of its inputs and develop systematic, abstract encodings of the environment. We found that comparison models based on standard error Bp learning did not learn a categorical structure that went beyond the surface similarity present in the visual input layers, and future work is focused on narrowing down the specific mechanisms required to drive this learning.

In addition to the predictive learning functions of the deep/thalamic layers, these same circuits are also likely critical for supporting powerful top-down attentional mechanisms that have a net multiplicative effect on superficial-layer activations (Bortone, Olsen, & Scanziani, 2014; Olsen, Bortone, Adesnik, & Scanziani, 2012). The importance of the pulvinar for attentional processing has been widely documented (e.g., Saalmann et al., 2012; Bender & Youakim, 2001; LaBerge & Buchsbaum, 1990), and there is likely an additional important role of the thalamic reticular nucleus, which can contribute a surround-inhibition contrast-enhancing effect on top of the incoming attentional signal from the cortex (Jaramillo, Mejias, & Wang, 2019; Wimmer et al., 2015; Pinault, 2004; Crick, 1984). In other work in progress, we have shown that the deep/thalamic circuits in our model produce attentional effects consistent with the abstract Reynolds and Heeger (2009) model, whereas the contributions of the deep layer networks to this function are broadly consistent with the folded-feedback model (Grossberg, 1999). These attentional modulation signals cause the bidirectional constraint satisfaction process in the superficial network to focus on task-relevant information while down-regulating responses to irrelevant information—in the real world, there are typically too many objects to track at any given time, so predictive learning must be directed toward the most important objects (Richter & de Lange, 2019; Cavanagh et al., 2010; Pylyshyn, 1989).

There are also data suggesting that the pulvinar is important for supporting confidence judgments, driven by relative ambiguity in a random dot motion categorization task (Komura et al., 2013). Critically for the present framework, this confidence modulation only emerged in the period after the first 100 msec of processing and manifested as a positive correlation with confidence (i.e., more unambiguous stimuli resulted in higher firing rates). We can interpret this as reflecting an ongoing generative postdiction of the stimulus signal, with stronger firing associated with more unambiguous top-down activation based on the current internal representation. Note that this directionality is the opposite of EE coding neurons, which would presumably increase with increasing error/ambiguity in the prediction. Interestingly, inactivation of these pulvinar neurons resulted in a substantial (200%) increase in opt-out choices on the most ambiguous stimuli, suggesting a

level of metacognitive awareness of the pulvinar signal (or at least a direct effect of pulvinar on relevant metacognitive processes). Predictive accuracy would be an ideal source of metacognitive confidence signals across a wide range of domains, suggesting another important contribution of pulvinar even after initial learning. Jaramillo et al. (2019) present a comprehensive model of attentional, decision-making, and working memory contributions of the pulvinar, including these confidence data, which is generally compatible with our framework, although it does not address any learning phenomena.

There are a number of important limitations of the current What–Where Integration (WWI) model, in terms of its scale and ability to process real-world cluttered visual scenes with multiple objects present, such as those used in the widely studied ImageNet data set. The model is much smaller than standard DCNN vision models, because its computational demands are significantly higher, in a way that also does not fit well with current graphics processing unit (GPU)-based parallel computation hardware, because of the relative complexity of the algorithms and the sparseness of the activations. For each image, 100 cycles (of 1 msec each) of activation updating are required to enable the bidirectional activation and inhibition to integrate in a graded manner over the alpha cycle, compared to only one such iteration for most feedforward DCNN models. Furthermore, the bidirectional connectivity, extensive shortcut connections, and use of multiple cortical lamina per cortical area result in significant increases in the number of synaptic connections, which dominate the computational cost, and scale roughly as n^2 in the number of neurons n per layer across one projection. Thus, there are 207 million connections for the full WWI model, requiring 10 GB of RAM, and it takes over a day to run using 32 high-performance CPU processors with fast network interconnects, using the fastest combination of threading and parallel batch training. Doubling the network size causes it to no longer fit in available RAM, and yet, its high-resolution V1 layer is only 16×16 , compared to 55×55 for basic DCNN models such as AlexNet and 224×224 for VGG16. The result is that the model has a relatively low-resolution view of the world, as reflected in the reconstructed images shown in Figure 5.

In addition to having a higher-resolution input to be able to process more complex real-world cluttered images, the model would require functional attentional dynamics to focus processing on a small number of objects at a time, as is well documented for humans processing complex images. Thus, once the attentional dynamics are well integrated with the predictive learning mechanisms, we can begin to explore performance on more complex images, subject to improved computational hardware supporting larger network sizes.

Considerable further work remains to be done to more precisely characterize the essential properties of our biologically motivated model necessary to produce this abstract form of learning and to further explore the full scope of predictive learning across different domains.

We strongly suspect that extensive cross-modal predictive learning in real-world environments, including between sensory and motor systems, is a significant factor in infant development and could greatly multiply the opportunities for the formation of higher-order abstract representations that more compactly and systematically capture the structure of the world (Yu & Smith, 2012). Future versions of these models could thus potentially provide novel insights into the fundamental question of how deep an understanding a preverbal human, or a nonverbal primate, can develop (Elman et al., 1996; Spelke, Breinlinger, Macomber, & Jacobson, 1992), based on predictive learning mechanisms. This would then represent the foundation upon which language and cultural learning builds, to shape the full extent of human intelligence.

APPENDIX

All of the materials described here, including the experimental study, the computational models, and the code to perform the representational similarity analysis, are all available on our github account at github.com/ccnlab/deep-obj-cat, and the new version of the emergent simulation environment is at github.com/emeral/leabra, which contains extensive documentation and examples that can be run in Python or the Go language. The best place to start in understanding computationally how the predictive learning model works is with the FSA model described in the main text, which is available at github.com/emeral/leabra/tree/master/examples/deep_fsa. For the large and complex WWI model, the most complete understanding can only be had by directly examining the code, as there are a number of details that are not efficiently captured in this Appendix text.

REPRESENTATIONAL SIMILARITY ANALYSIS METHODS

The different representations being compared here are the following:

Leabra: The DeepLeabra (biological model) TE layer representations (specifically TEs = superficial—results are very similar for deep as well).

Bp: The TEs layer representations from the Bp version of biological model, including “what,” “where,” and “What \times Where” integration layers, trained with the V1p and V1hp (low- and high-resolution pulvinar) layers as the final output layers, using the time t target pattern from the $t - 1$ input (i.e., as a predictive network).

V1: The Gabor-filtered representation of the visual input to both of the above models, which was identical across them.

PredNet: The highest layer (sixth layer) of the PredNet architecture.

Expt: Similarity matrix constructed from human pairwise similarity judgments (see Behavioral Experiment Methods).

An optimal category cluster can be defined as one that has high within-cluster similarity and low between-cluster similarity. This can be operationalized by the contrast distance metric, based on a 1-correlation (dissimilarity) measure, as the difference between the average within-cluster similarity and the average between-cluster similarity:

$$CD = (1 - r_{in}) - (1 - r_{out}) \quad (1)$$

With distance-like 1-correlation values, this contrast distance should be minimized (it is typically negative), or equivalently, the contrast on raw correlation values can be maximized (it is typically a positive number—just the sign flip of distance value). We refer to the positive numbers and maximization here as that is more intuitive.

Starting with an initial set of clusters, a permutation-based hill-climbing strategy was used to determine a local minimum in this measure: Each item was tested in each of the other possible categories, and if that configuration reduced the overall average contrast distance (ACD) metric across all items, then it was adopted and the process iterated until no such permutation improved the metric. This algorithm can only decrease the number of clusters (by moving all items out of a given cluster), so different numbers of initial clusters can be used to search the overall space.

Figure 14 shows the resulting categories. The Bp model converged on the same cluster state from all starting configurations tested, varying from five to two initial categories. This is the cluster set shown in Figure 10 of the

main paper and has an ACD of 0.0838 (this is relatively low because the patterns were overall quite similar). Likewise, the V1 patterns (which were the same across Leabra and Bp models) reliably converged on the same pattern (shown in Figure 10), with ACD = 0.2448.

For the PredNet Layer 6 representations, starting from the V1 categories gave the best results of any other set (ACD = 0.1967), and a few permutations resulted in a reliable solution that was arrived at from all other three category starting points tested, shown in Figure 14 (ACD = 0.2820). This indicates that PredNet did not go much beyond the structure present in the input, although it did not use the V1 Gabor filtering used in the Leabra and Bp models (i.e., this V1-level encoding well captures the structure of the visual inputs in general). The PredNet pixel and Layer 1 representations both converged on essentially a single monolithic category with very low ACDs (0.0018 and 0.0013, respectively).

For the Leabra TE representations, we found a set of centroid-shape categories that are near-best when considering both the Leabra model and the results from the human behavioral experiment. Starting from these categories, the permutation analysis converged on reducing the size of the vertical and round categories to one item each, over a sequence of five steps. This is consistent with the observation from Figure 7 that there are three broader categories within which the five finer-grained categories are embedded (i.e., vertical and pyramid are overall similar to each other, as are round and box). Nevertheless, our

Figure 14. Shape categories used for similarity matrix plots in the main paper. “Centroid” shape categories are near-best for both the Leabra model and the Expt results and fit our visual intuitions about overall shape. “Bp” is reliably optimal for the Bp model from all starting points. “V1” reliably optimal for V1 inputs and also was close to the best for the Bp and PredNet Layer 6 representations. “PredNet” is the best stable solution for PredNet Layer 6.

Centroid		Bp	
1. pyramid	3. round cont'd	1. cat1	1. cat1 cont'd
• banana	• handgun	• banana	• handgun
• layercake	• chair	• layercake	• chair
• trafficcone	4. box	• trafficcone	• slrcamera
• sailboat	• slrcamera	• sailboat	• elephant
• trex	• elephant	• trex	• piano
2. vertical	• piano	• person	• fish
• person	5. horiz	• guitar	• car
• guitar	• fish	• tablelamp	2. cat2
• tablelamp	• car	• doorknob	• heavycannon
3. round	• heavycannon	• donut	• stapler
• doorknob	• stapler		• motorcycle
• donut	• motorcycle		
V1		PredNet	
1. cat1	2. cat2 cont'd	1. cat1	2. cat2 cont'd
• trafficcone	• handgun	• trafficcone	• slrcamera
• sailboat	• slrcamera	• sailboat	• elephant
• person	• elephant	• person	• fish
• guitar	• piano	• guitar	• car
• tablelamp	• fish	• tablelamp	• heavycannon
• chair	• car	• layercake	• stapler
2. cat2	• heavycannon	2. cat2	• motorcycle
• layercake	• stapler	• trex	3. cat3
• trex	• motorcycle	• donut	• chair
• doorknob	3. cat3	• banana	• doorknob
• donut	• banana	• handgun	• piano

initial visual intuition about the broad shape categories, along with a bias against having single-item categories, reinforced the use of the finer-grained centroid selection. The average contrast difference of our centroid selection is 0.5071, whereas the maximal result from the permutation was 0.5526, which is a relatively small proportional difference.

Furthermore, once we had collected the human experimental data (Expt), it was clear that it strongly coincided with our original shape intuitions and with the finer-grained five-category centroid structure. Starting from the centroid categories, the maximal permutation made only three changes, moving trex (T-rex) and handgun into the horizontal category, and chair into the pyramid, going from a distance score of 0.3083 to 0.3225, which is a relatively small improvement. However, using the maximal Expt clusters directly on the Leabra model gives a lower ACD measure of 0.3745 (compared to 0.5071 for centroid), so the centroid categories represent a good middle ground between Expt and the model, and this strong shared similarity structure with near-optimal cluster structures confirms that the model and people are encoding largely the same information.

In contrast, if we organize the Expt similarity matrix using the Bp categories, it produces a very poor ACD measure of 0.0643 (compared to 0.3083 for the centroid categories), strongly suggesting that people's shape representations are not compatible with that simple structure.

Another approach to determining clusters from similarity matrices, "agglomerative clustering," starts with all items as singletons and iteratively combines the closest two into a new cluster. The results for the Leabra and Expt similarity matrices are shown in Figure 15, which has also color-coded the items in terms of their category status according to the centroid structure. Because of a strong history dependency in the clustering process and the indeterminacy of reducing a high-dimensional similarity structure down to two dimensions, structure beyond the leaf level is not very reliable (ties are also broken by a random number generator), but nevertheless, you can clearly see that, in both cases, items from the same cluster are almost always together as leaves in the plots. This then provides additional converging support for the idea that the model is learning the same kind of shape categories as people have.

For the network layer RSA computations, activation vectors were accumulated separately for each 3-D object item and, within that, separately for each frame index of the movie. To be able to monitor similarity metrics as the model trained, we used a running-average integration of neural activity across trials to accumulate the patterns. Specifically, the current activation pattern across each layer was recorded and averaged unit-by-unit with a time constant of $\tau = 10$. Critically, by integrating separately for each frame, this running-average computation did not introduce any bias for temporally adjacent frames to be more similar. Nevertheless, when we computed the frame-to-frame

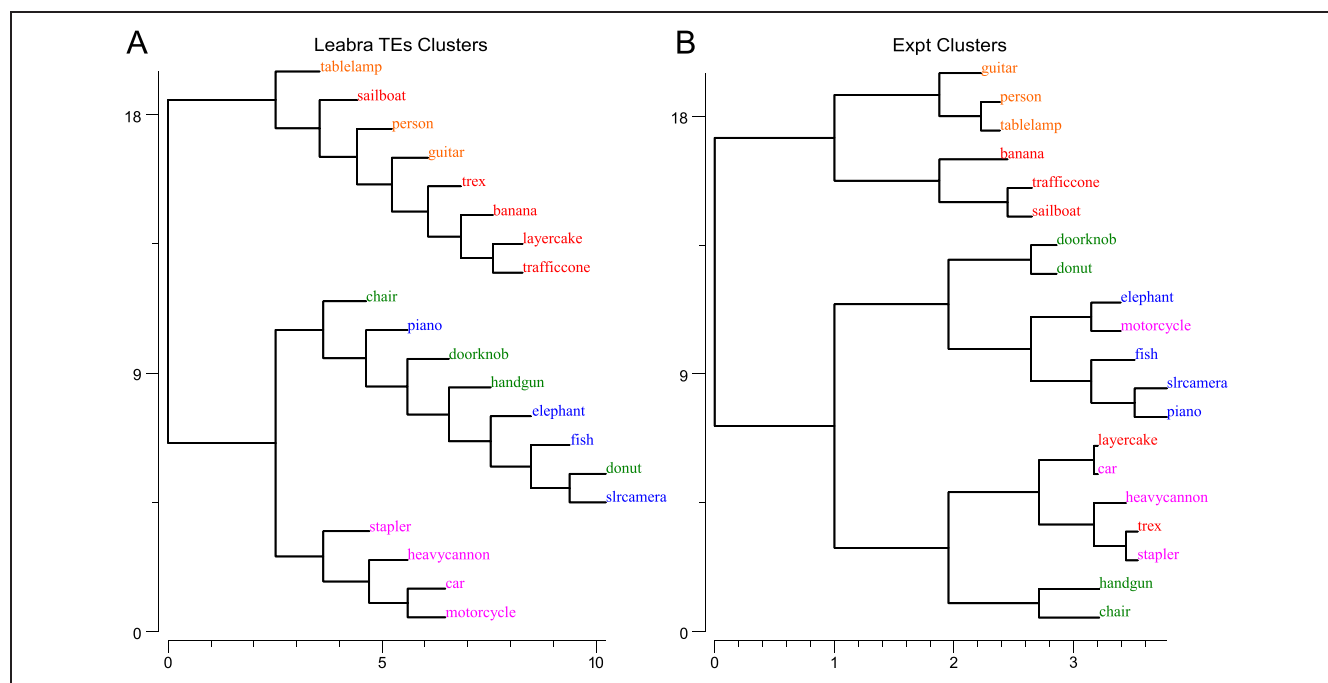


Figure 15. Agglomerative clustering on the Leabra and Expt representations, with the centroid categories color coded. The most reliable information from this is the leaf-level groupings, as the rest of the structure is indeterminate and history dependent in reducing higher-dimensional structure down to a 2-D plot. Both cluster plots show a strong tendency to group leaf items together in the same centroid categories, with a few exceptions in each case. In addition, the Leabra plot nicely captures the broader three-category structure evident in the similarity matrix plots, within which the five finer-grained centroid categories are organized. Overall, this provides further confirmation that the model and the human participants are organizing the shapes in largely the same way.

similarities for TE, they were quite high (.901 correlation on average across all objects).

BEHAVIORAL EXPERIMENT METHODS

The behavioral experiment was conducted on Amazon.com's MTurk Web platform under University of Colorado institutional review board approval (19-0176), using 30 participants each categorizing up to 800 image pairs as shown in Figure 16, using the standard simple image categorization framework with a lightly customized script. Objects were drawn from the 156 3-D object set, but data were aggregated in terms of the 20 basic-level categories (car, stapler, etc.) because we could not sample all 156×156 object pairs. Thus, the resulting data were aggregated for each category pair in terms of the proportion of times when that pair was selected when presented.

The individual images were produced by reconstructing from the V1 transform that the computational model used in its high-resolution V1 input layer, to give human participants as similar of an experience as possible to how the model "saw" the objects, and to reduce the influence of existing semantic knowledge, which was entirely missing in our model (Figure 16).


BIOLOGICAL MODEL METHODS

This section provides more information about the DeepLeabra WWI model. The purpose of this information is to give more detailed insight into the model's function beyond the level provided in the main text, but with a model of this complexity, the only way to really understand it is to explore the model itself. It is available for download at github.com/ccnlab/deep-obj-cat/tree/master/sims/cemer. We now have a full replication of this model in our new, much more transparent simulation framework, available at github.com/ccnlab/deep-obj-cat/tree/master/sims/wwi3d—this is more readable and recommended. Furthermore, the best way to understand this model is to understand the framework in which it is implemented, which is explained in great detail, with many running simulations explaining specific elements of functionality, at CompCogNeuro.org.

Layer Sizes and Structure

Figure 5 in the main text shows the general configuration of the model, and Table 1 shows the specific sizes of each of the layers and where they receive inputs from.

Figure 16. Example stimulus from the behavioral experiment, using the V1 reconstruction of the actual input images presented to the model, to better capture the coarse-grained perception of the model. Participants were requested to choose which of the two pairs, left or right, was most similar in terms of overall shape.



All the activation and general learning parameters in the model are at their standard Leabra defaults.

Projections

The general principles and patterns of connectivity are shown in Figure 17 (and Figures 1 and 2 in the main text). As noted in the main text, the connectivity and overall structure obeys the established principles identified in neocortical anatomy (Markov, Ercsey-Ravasz, et al., 2014; Markov, Vezoli, et al., 2014; Felleman & Van Essen, 1991; Rockland & Pandya, 1979).

Detailing each of the specific parameters associated with the different projections shown in Table 1 would take too much space—those interested in this level of detail should download the model from the link shown above. There are topographic projections between many of the lower-level retinotopically mapped layers, consistent with our earlier vision models (O'Reilly et al., 2013). For example the 8×8 unit groups in V2 are reduced down to the 4×4 groups in V3 via a 4×4 unit-group topographic projection, where neighboring units have half-overlapping receptive fields (i.e., the field moves over two unit groups in V2 for every one unit group in V3), and the full space is uniformly tiled by using a wraparound effect at the edges. Similar patterns of connectivity are used in standard DCNNs. However, we do not share weights across units as in a true convolutional network.

The projections from ObjVel (object velocity) and SaccadePlan layers to LIPs, LIPd were initialized with a topographic sigmoidal pattern that moved as a function of the position of the unit group, by a factor of .5, whereas the projections from EyePos were initialized with a Gaussian pattern. These patterns multiplied uniformly distributed random weights in the .25–.75 range, with the lowest values in the topographic pattern having a multiplier of .6, whereas the highest had a multiplier of 1 (i.e., a fairly subtle effect). This produced faster convergence of the LIP layer when doing "where" pathway pretraining compared to purely random initial weights, consistent with Pouget and Sejnowski (1997) and related work on parietal gain field basis function representations.

In addition to exploring different patterns of overall connectivity, we also explored differences in the relative strengths of receiving projections, which can be set with a `wt_scale.rel` parameter in the simulator. All feedforward pathways have a default strength of 1. For the feedback

Table 1. Layer Sizes, Showing Numbers of Units in One Pool (or Entire Layer if Pool is Missing), and the Number of Pools of Such Units, along x, y Axes

Area	Name	Units		Pools		Receiving Projections
		x	y	x	y	
V1	V1s	4	5	8	8	
	V1p	4	5	8	8	V1s, V2d, V3d, V4d, TEOd
V1h	V1hs	4	5	16	16	
	V1hp	4	5	16	16	V1s, V2d, V3d, V4d, TEOd
Eyes	EyePos	21	21			
	SaccadePlan	11	11			
	Saccade	11	11			
Obj	ObjVel	11	11			
V2	V2s	10	10	8	8	V1s, LIPs, V3s, V4s, TEOd, V1p, V1hp
	V2d	10	10	8	8	V2s , V1p, V1hp, LIPd, LIPp, V3d, V4d, V3s, TEOs
LIP	MtPos	1	1	8	8	V1s
	LIPs	4	4	8	8	MtPos, ObjVel, SaccadePlan, EyePos, LIPp
	LIPd	4	4	8	8	LIPs , LIPp, ObjVel, Saccade, EyePos
	LIPp	1	1	8	8	MtPos , V1s, LIPd
V3	V3s	10	10	4	4	V2s, V4s, TEOs, DPs, LIPs, V1p, V1hp, DPp, TEOd
	V3d	10	10	4	4	V3s , V1p, V1hp, DPp, LIPd, DPd, V4d, V4s, DPs, TEOs
	V3p	10	10	4	4	V3s , V2d, DPd, TEOd
DP	DPs	10	10			V2s, V3s, TEOs, V1p, V1hp, V3p, TEOp
	DPd	10	10			DPs , V1p, V1hp, DPp, TEOd
	DPp	10	10			DPs , V2d, V3d, DPd, TEOd
V4	V4s	10	10	4	4	V2s, TEOs, V1p, V1hp
	V4d	10	10	4	4	V4s , V1p, V1hp, V4p, TEOd, TEOs
	V4p	10	10	4	4	V4s , V2d, V3d, V4d, TEOd

Downloaded from http://direct.mit.edu/jocn/article-pdf/33/6/1158/1913577/jocn_a_01708.pdf by guest on 07 September 2023

Table 1. (continued)

Area	Name	Unites		Pools		Receiving Projections
		<i>x</i>	<i>y</i>	<i>x</i>	<i>y</i>	
TEO	TEOs	10	10	4	4	V4s, V1p, V1hp, TEs
	TEOd	10	10	4	4	TEOs, TEOd , V1p, V1hp, V4p, TEOp, TEp, TEd
	TEOp	10	10	4	4	TEOs , V3d, V4d, TEOd, TEd
TE	TEs	10	10	4	4	TEOs, V1p, V1hp
	TEd	10	10	4	4	TEs, TEd , V1p, V1hp, V4p, TEOp, TEp, TEOd
	TEp	10	10	4	4	TEs , V3d, V4d, TEOd

Each area has three associated layers: s = superficial layer; d = deep layer (context updated by 51B neurons in the same area, shown in **bold**); and p = pulvinar layer (driven by 51B neurons from the associated area, shown in **bold**).

projections, which are typically weaker (consistent with the biology), we explored a discrete range of strengths, typically .5, .2, .1, and .05. The strongest top-down projections were into V2s from LIP and V3, whereas most others were .2 or .1. Likewise, projections from the pulvinar were weaker, typically .1. These differences in strength sometimes had large effects on performance during the initial bootstrapping of the overall model structure, but in the final model, they are typically not very consequential for any individual projection.

Training Parameters

Training typically consisted of 512 alpha trials per epoch (51.2 sec of real-time equivalent), for 1000 such epochs. Each trial was generated from a virtual reality environment in the emergent simulator, which rendered first-person views with moving eye position onto the object tumbling through space with fixed motion and rotation parameters over the sequence of eight frames (see Figure 5 in the main text for a representative example). Each frame was rendered at a 256 × 256 resolution and processed through our standard V1 Gabor filters, which are described in detail in O'Reilly et al. (2013).

Because the start of each sequence of eight frames is unpredictable, we turned off learning for that trial, which improves learning overall. We have recently developed an automatic such mechanism based on the running average (and running variance) of the prediction error, where we turn off learning whenever the current prediction error *z*-normalized by these running average values is below 1.5 *SDs*, which works well and will be incorporated into future models. Biologically, this could correspond to a connection between pulvinar and neuromodulatory areas that could regulate the effective learning rate in this way.

Figure 18A shows the learning trajectory of the model, indicating that it learns quite rapidly. This rapid initial learning is likely facilitated by the extensive use of shortcut

connections converging from all over the simulated visual system onto the V1 pulvinar layers and direct projections back from these pulvinar layers. Thus, error signals are directly communicated and can drive learning quickly and efficiently. However, there are also extensive indirect, bi-directional connections among the superficial layers, which can drive indirect error Bp learning as well.

Model Algorithms

The biologically based model was implemented using the Leabra framework, which is described in detail in previous publications (O'Reilly et al., 2012, 2016; O'Reilly & Munakata, 2000; O'Reilly, 1996, 1998), and summarized here. The online textbook at CompCogNeuro.org provides the most comprehensive description of the framework, and github.com/emergent/leabra has a summary of all the equations (and the code itself). There are two main implementations of Leabra, one in the C++ emergent software and a new one using Go and Python language at the prior link. These same equations and standard parameters have been used to simulate over 40 different models in O'Reilly and Munakata (2000), O'Reilly et al. (2012), and a number of other research models. Thus, the model can be viewed as an instantiation of a systematic modeling framework using standardized mechanisms, instead of constructing new mechanisms for each model (O'Reilly et al., 2016).

The neurons use a rate code version of the adaptive exponential conductance-based point neuron model (Brette & Gerstner, 2005), with the standard resistor-capacitor circuit equations:

$$\Delta V_m(t) = \tau \sum_c g_c(t) \bar{g}_c (E_c - V_m(t)), \quad (2)$$

where *c* represents excitatory, inhibitory, and leak channels. Inhibition is driven by simulated interneurons in proportion

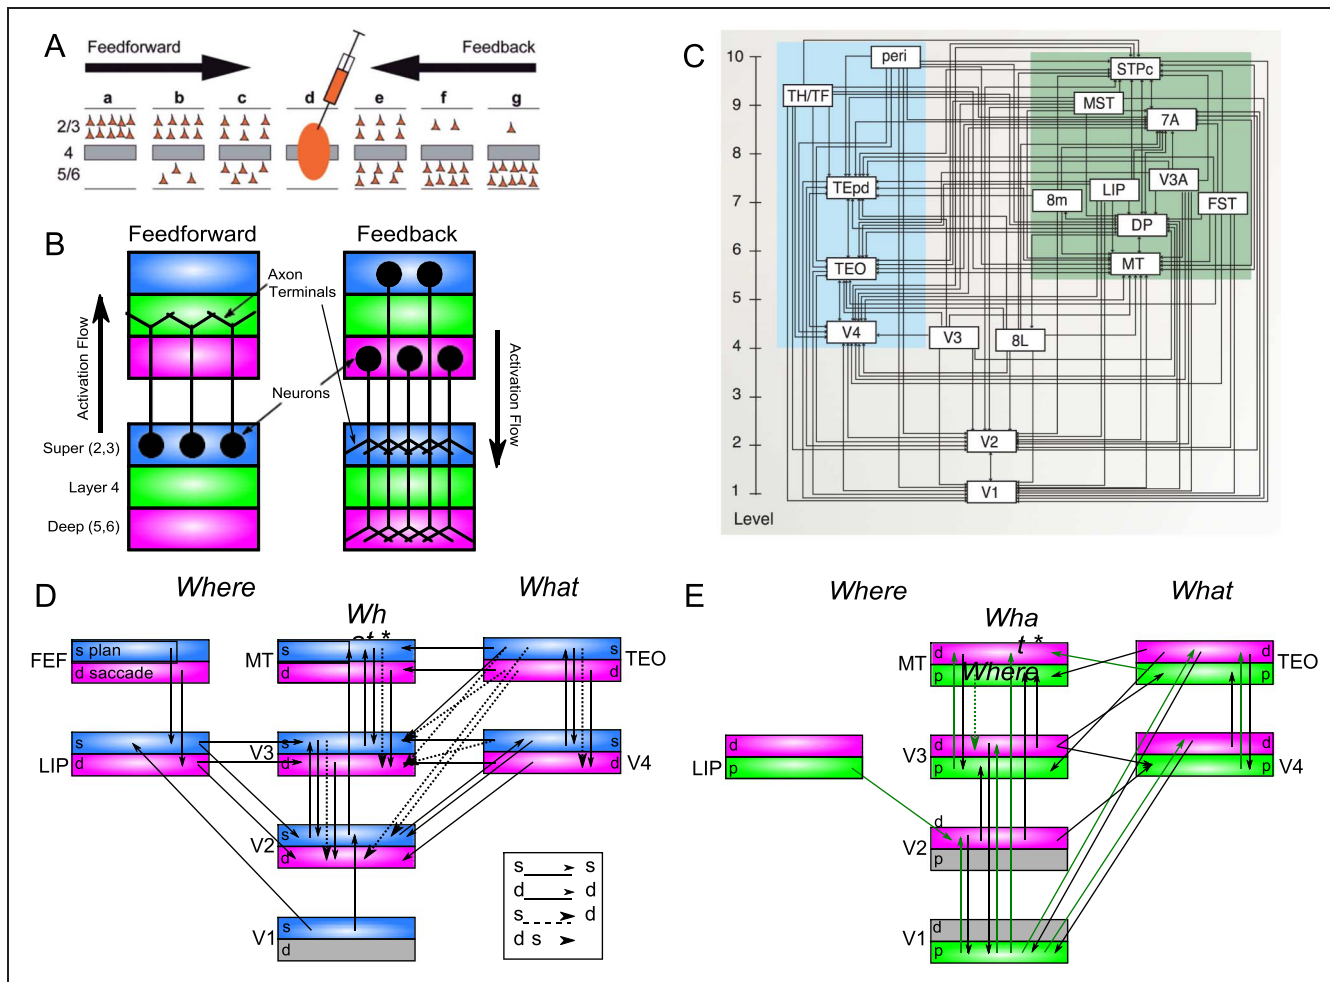


Figure 17. Principles of connectivity in DeepLeabra. (A) Markov et al. (2014) data showing density of retrograde labeling from a given injection in a middle-level area (“d”): Most feedforward projections originate from superficial layers of lower areas (“a–c”), and deep layers predominantly contribute to feedback (and more strongly for longer-range feedback). (B) Summary diagram showing most feedforward connections originating in superficial layers of lower area and terminating in Layer 4 of higher areas, whereas feedback connections can originate in either superficial or deep layers and, in both cases, terminate in both superficial and deep layers of the lower area (adapted from Felleman & Van Essen, 1991). (C) Anatomical hierarchy as determined by percentage of superficial layer source labeling by Markov et al. (2014)—the hierarchical levels are well matched for our model, but we functionally divide the dorsal pathway (shown in green background) into the two separable components of a “where” pathway and a What × Where integration pathway. (D) Superficial and deep-layer connectivity in the model. Note the repeating motif between hierarchically adjacent areas, with bidirectional connectivity between superficial layers, and feedback into deep layers from both higher-level superficial and deep layers, according to the canonical pattern shown in A and B. Special patterns of connectivity from TEO to V3 and V2, involving crossed super-to-deep and deep-to-super pathways, provide top-down support for predictions based on high-level object representations. (E) Connectivity for deep layers and pulvinal in the model, which generally mirror the corticocortical pathways (in D). Each pulvinal layer (p) receives 5IB driving inputs from the labeled layer (e.g., V1p receives 5IB drivers from V1). In reality, these neurons are more distributed throughout the pulvinal, but it is computationally convenient to organize them together as shown. Deep layers (“d”) provide predictive input into the pulvinal, and pulvinal projections send error signals (via temporal differences between predictions and actual state) to both deep and superficial layers of given areas (only “d” shown). Most areas send deep-layer prediction inputs into the main V1p prediction layer and receive reciprocal error signals therefrom. The strongest constraint we found was that pulvinal outputs (colored green) must generally project only to higher areas, not to lower areas, with the exceptions of DPP → V3 and LIPp → V2. V2p was omitted because it is largely redundant with V1p in this simple model.

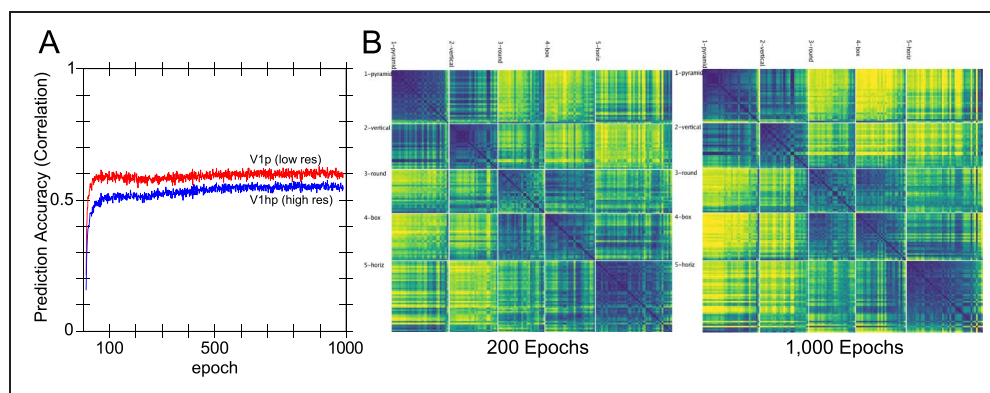
to feedforward and feedback dynamics, producing sparse distributed representations and controlling the effects of bidirectional excitatory connections between layers.

Each neuron learns using a more biologically based version of the contrastive Hebbian learning (CHL) algorithm, as shown in Figure 2:

$$\Delta_{\text{CHL}} = x^+y^+ - x^-y^- \quad (3)$$

where x is the sending activation, y is the receiving activation, and the $+$ and $-$ superscripts indicate activations in the plus and minus phases, respectively. The actual learning equations, detailed at github.com/emery/leabra and in the online textbook at CompCogNeuro.org, produce a combination of error-driven and self-organizing factors, which emerge out of a single learning rule that was derived from a biologically detailed model of synaptic plasticity by

Figure 18. (A) Predictive learning curve for DeepLeabra, showing the correlation between prediction and actual over the two different V1 layers. Initial learning is quite rapid, followed by a slower but progressive learning process that reflects development of the IT representations (e.g., manipulations that interfere with those areas selectively impair this part of the learning curve). Overall prediction accuracy remains far from perfect, as shown in Figure 5 in the main text, and significantly worse than the Bp-based models. This is a typical finding from Leabra models, which are significantly more constrained as a result of bidirectional attractor dynamics, Hebbian learning, and inhibitory competition—that is, the very things that are likely important for forming abstract categorical representations. (B) Similarity matrix over TEs layer at 200 epochs, which has less contrast and definition (particularly evident in the off-block-diagonal differences) compared to the 1000-epoch result (C; also shown in Figure 7 in the main text).



(Urakubo et al., 2008) and is closely related to the BCM algorithm (Bienenstock et al., 1982).

Deep Context

This section describes in detail the equations that are specific to the “deep” version of Leabra that implements the specific predictive learning additions to the general algorithm. Like the SRN (Elman, 1990; Jordan, 1989), which the deep predictive learning model functionally resembles, the primary computational specialization required is the maintenance of prior temporal context in the CT layer. In addition, the pulvinar layers have to be driven by the bottom-up inputs in the plus phase, after being driven by the CT inputs in the minus phase.

Computationally, the CT layer is specialized for maintaining context from the previous alpha cycle, to generate the prediction over the pulvinar layer. At the end of every plus phase, a new CT context excitatory input is computed from the normalized dot product of the context weights times the sending activations, just as in the standard net input used in Leabra:

$$\eta_j = \langle x_i w_{ij} \rangle = \frac{1}{n} \sum_i x_i w_{ij} \quad (4)$$

where x_i are the sending activations and w_{ij} are the weights. This net input is then added in with the standard net input at each cycle of processing during the subsequent alpha cycle.

The relative strength of these context layer inputs was set progressively larger for higher layers in the network, with a maximum of four in V4, TEO, and TE. In addition, TEO and TE received “self” context projections, which provide an extended window of temporal context into the prior 200-msec interval, consistent with multiple sources of neural data (Chaudhuri et al., 2015). These self projections were connected only within the narrower pool level of units, enabling these neurons to develop mutually excitatory loops to sustain activations over the multiple

trials when the same object was present. We hypothesize that these modifications correspond to biological adaptations in IT cortex that likewise support greater sustained activation of object-level representations.

Learning of the context weights occurs as normal, but using the sending activation states from the prior time step’s activation.

Computational and Biological Details of SRN-like Functionality

Predictive autoencoder learning has been explored in various frameworks, but the most relevant to our model comes from the application of the SRN to a range of predictive learning domains (Elman et al., 1996; Elman, 1990). One of the most powerful features of the SRN is that it enables error-driven learning, instead of arbitrary parameter settings, to determine how prior information is integrated with new information. Thus, SRNs can learn to hold onto some important information for a relatively long interval, while rapidly updating other information that is only relevant for a shorter duration. This same flexibility is present in our DeepLeabra model. Furthermore, because this temporal context information is hypothesized to be present in the deep layers throughout the entire neocortex (in every microcolumn of tissue), the DeepLeabra model provides a more pervasive and interconnected form of temporal integration compared to the SRN, which typically only has a single temporal context layer associated with the internal “hidden” layer of processing units.

An extensive computational analysis of what makes the SRN work as well as it does, and explorations of a range of possible alternative frameworks, has led us to an important general principle: Subsequent outcomes determine what is relevant from the past. At some level, this may seem obvious, but it has significant implications for predictive learning mechanisms based on temporal context. It means that the information encoded in a temporal context




Figure 19. How the DeepLeabra temporal context computation compares to the SRN mathematically. (A) In a standard SRN, the context (deep layer biologically) is a copy of the hidden activations from the prior time step, and these are held constant while the hidden layer (superficial) units integrate the context through learned synaptic weights. (B) In DeepLeabra, the deep layer performs the weighted integration of the soon-to-be context information from the superficial layer and then holds this integrated value and feeds it back as an additive net-input-like signal to the superficial layer. The context net input is precomputed, instead of having to compute this same value over and over again. This is more efficient and more compatible with the diffuse interconnections among the deep layer neurons. Layer 6 projections to the thalamus and back recirculate this precomputed net input value into the superficial layers (via Layer 4) and back into itself to support maintenance of the held value.

representation cannot be learned at the time when that information is presently active. Instead, the relevant contextual information is learned on the basis of what happens next.

This explains the peculiar power of the otherwise strange property of the SRN: The temporal context information is preserved as a “direct copy” of the state of the hidden layer units on the previous time step (Figure 19), and then learned synaptic weights integrate that copied context information into the next hidden state (which is then copied to the context again, and so on). This enables the error-driven learning taking place in the current time step to determine how context information from the previous time step is integrated. Furthermore, the simple direct copy operation eschews any attempt to shape this temporal context itself, instead relying on the learning pressure that shapes the hidden layer representations to also shape the context representations. In other words, this copy operation is essential, because there is no other viable source of learning signals to shape the nature of the context representation itself (because these learning signals require future outcomes, which are by definition only available later).

The direct copy operation of the SRN is however seemingly problematic from a biological perspective: How could neurons copy activations from another set of neurons at some discrete point in time and then hold onto those copied values for a duration of 100 msec, which is a reasonably

long period in neural terms (e.g., a rapidly firing cortical neuron fires at around 100 Hz, meaning that it will fire 10 times within that context frame)? However, there is an important transformation of the SRN context computation, which is more biologically plausible and compatible with the structure of the deep network (Figure 19). Specifically, instead of copying an entire set of activation states, the context activations (generated by the phasic 5IB burst) are immediately sent through the adaptive synaptic weights that integrate this information, which we think occurs in the 6CC (corticocortical) and other lateral integrative connections from 5IB neurons into the rest of the deep network.

The result is a precomputed net input from the context onto a given hidden unit (in the original SRN terminology), not the raw context information itself. Computationally, and metabolically, this is a much more efficient mechanism, because the context is, by definition, unchanging over the 100-msec alpha cycle, and thus, it makes more sense to precompute the synaptic integration, rather than repeatedly recomputing this same synaptic integration over and over again (in the original feedforward Bp-based SRN model, this issue did not arise because a single step of activation updating took place for each context update—whereas in our bidirectional model, many activation update steps must take place per context update).

There are a couple of remaining challenges for this transformation of the SRN. First, the precomputed net input from the context must somehow persist over the subsequent 100-msec period of the alpha cycle. We hypothesize that this can occur via *N*-methyl-D-aspartate and metabotropic glutamate receptor channels that can easily produce sustained excitatory currents over this time frame. Furthermore, the reciprocal excitatory connectivity from 6CT to TRC and back to 6CT could help to sustain the initial temporal context signal. Second, these contextual integration synapses require a different form of learning algorithm that uses the sending activation from the prior 100 msec, which is well within the time constants in the relevant calcium and second messenger pathways involved in synaptic plasticity.

BACKPROPAGATION MODEL METHODS

The Bp version of the WWI model has the same layer sizes and feedforward patterns of connectivity as the DeepLeabra version. Topographically, the V1p and V1hp pulvinar layers serve as output layers at the highest level of the network, receiving all the various connections from deep layers as shown in Table 1. Likewise, the LIPp served as a target output layer for the “where” pathway. To achieve predictive learning, the V1 pulvinar targets were from the scene at time t , whereas the V1s inputs were from the scene at time $t - 1$. We also ran a comparison autoencoder model that had inputs and target outputs from the same time step, and it showed even less systematic organization of its higher-level representations, further supporting the notion that predictive learning is important, across all frameworks.

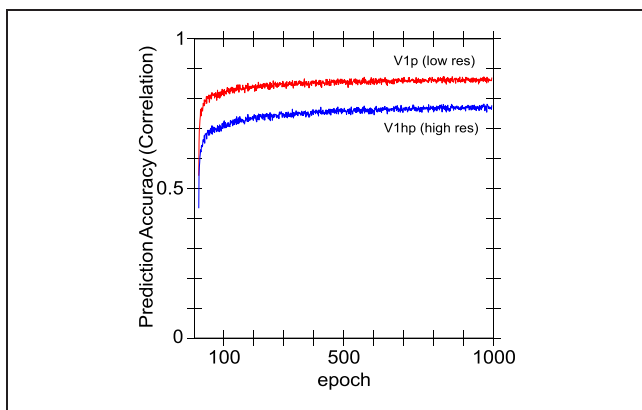


Figure 20. Learning curves for the Bp version of the WWI model. Although it achieves better predictive accuracy than the DeepLeabra version, it fails to acquire abstract object category structure, indicating a potential tradeoff between simplifying and categorizing inputs, versus predicting precisely where the low-level visual features will move.

The learning curve for the predictive version is shown in Figure 20, which shows better overall prediction accuracy compared to the DeepLeabra model. However, as the RSA showed, this Bp model failed to learn object categories that go beyond the input similarity structure, indicating that perhaps it was paying too much “attention” in learning to this low-level structure, and lacked the necessary mechanisms to enable it to impose a simplifying higher-level structure on top of these inputs.

PREDNET MODEL METHODS

The PredNet architecture was designed to incorporate principles from predictive coding theory into a neural network model for predicting the next frame in a video sequence. Details of the model can be found in the original paper (Lotter et al., 2016), but here, we provide a brief overview of the architecture.

Architecture

PredNet is a DCNN that is composed of layers containing discrete modules. The lowest layer generates a prediction of incoming inputs (i.e., the pixels in the next frame), whereas each of the higher layers attempts to predict the errors made by the previous layer. Each layer contains an input convolutional module (A_l), a recurrent representational module (R_l), a prediction module (\hat{A}_l), and a representation of its own errors (E_l). The input convolutional module (A_l) transforms its input with a set of standard convolutional filters, a rectified linear activation function, and a max-pooling operation. The recurrent representation module (R_l) is a convolutional LSTM, which is a recurrent convolutional network that replaces the matrix multiplications in the standard long short-term memory (LSTM) equations with convolutions, allowing it to maintain a spatially organized representation of its inputs over time. The

prediction module (\hat{A}_l) consists of another standard convolutional layer and rectified linear activation that is used to generate predictions from the output of R_l . These predictions are then compared against the output of the input convolutional module (A_l). The errors generated in this comparison are represented explicitly in E_l , which applies a rectified linear activation to a concatenation of the positive ($A_l - \hat{A}_l$) and negative ($\hat{A}_l - A_l$) prediction errors. These errors then become the inputs to the next layer.

$$A_l^t = \begin{cases} x_t & \text{if } l = 0 \\ \text{MaxPool}(\text{ReLU}(\text{Conv}(E_{l-1}^t))), & \text{if } l > 0 \end{cases} \quad (5)$$

$$\hat{A}_l^t = \text{ReLU}(\text{Conv}(R_l^t)) \quad (6)$$

$$E_l^t = [\text{ReLU}(A_l^t - \hat{A}_l^t); \text{ReLU}(\hat{A}_l^t - A_l^t)] \quad (7)$$

$$R_l^t = \text{ConvLSTM}(E_l^{t-1}, R_l^{t-1}, \text{UpSample}(R_{l+1}^t)) \quad (8)$$

At each time step in the video sequence, PredNet generates a prediction of the next frame. This is done as follows: First, the R_l is computed for each layer starting from the top of the hierarchy (because each R_l^t depends on input from R_{l+1}^t), and then the A_l^t , \hat{A}_l^t , and E_l^t are computed in a feedforward fashion (because each A_l^t depends on input from the layer below, E_{l-1}^t).

All analyses in the RSA were conducted using the representations from the R_l layers.

Implementation Details

All experiments with the PredNet architecture were performed using PyTorch. An informal hyperparameter search was conducted to find the settings that maximized representational similarity to the human judgments. This was done by conducting RSA on each layer for each hyperparameter setting and computing, according to the centroid categories derived from the human data, the difference between the average within-category similarity and the average between-category similarity. Our final




Figure 21. Learning curves for the PredNet model. This model achieves the best overall prediction performance but also has the least well-differentiated, categorical representations.

architecture had six layers with 3, 16, 32, 64, 128, and 256 filters in the A_i and R_i modules and 3×3 kernels throughout the whole network. We also found that using sigmoid and tanh activation functions in fully connected convolutional LSTMs slightly improved performance, so these were used for all experiments.

The weights in the PredNet model are trained using error Bp. Predictions are generated, and errors are computed at all levels of the hierarchy, but the model performs better when only the lowest layer's errors are backpropagated (Lotter et al., 2016). We confirmed these results with experiments that backpropagated the errors in higher layers, in which performance (in terms of mean squared error) was marginally reduced but the RSA results were similar. For this reason, all reported experiments used a PredNet that was trained by only backpropagating the lowest level error.

The model was trained using a batch size of 8 and an Adam optimizer with a learning rate of 0.0001, with no scheduler, for 150,000 batches. A training curve is shown in Figure 21, showing that it achieves the best overall prediction accuracy of any model we tested and yet does not have representations that are as differentiated or categorical as our biologically based model, as shown in the main paper.

Regularization Experiments

As discussed in the main paper, our biologically based model includes a number of important biologically motivated properties that may be contributing to the development of its categorical representations. These properties, including excitatory bidirectional connections, inhibitory competition, and an additional form of Hebbian learning, may be acting as regularizers that encourage categorical learning. We therefore tested whether standard regularization methods used in deep learning would have similar




Figure 22. Effect of dropout in PredNet on RSA, as measured by the difference between the average within-category correlation and the average between-category correlation (using the centroid categories derived from human data). Dropout marginally improves the category structure learned in PredNet.

effects on the representations developed in the PredNet architecture. We tested (1) batch normalization, (2) dropout (0.1, 0.3, and 0.5), and (3) weight decay (0.01, 0.001, 0.0001, 0.00001). All experiments with batch normalization and weight decay showed reduced performance (in terms of both prediction error on the test set and within-category correlation). As shown in Figure 22, dropout marginally improved the within-category correlation while also slightly improving prediction accuracy, so a dropout rate of 0.1 was used for the comparison to our biologically based model in the main paper.

Acknowledgments

We thank Dean Wyatte, Tom Hazy, Seth Herd, Kai Krueger, Tim Curran, David Sheinberg, Lew Harvey, Jessica Mollick, Will Chapman, Helene Devillez, and the rest of the CCN Lab for many helpful comments and suggestions. This work was supported by ONR grants ONR N00014-19-1-2684/N00014-18-1-2116, N00014-14-1-0670/N00014-16-1-2128, N00014-18-C-2067, N00014-13-1-0067, and D00014-12-C-0638.

This work utilized the Janus supercomputer, which is supported by the National Science Foundation (award number CNS-0821794) and the University of Colorado Boulder. The Janus supercomputer is a joint effort of the University of Colorado Boulder, the University of Colorado Denver, and the National Center for Atmospheric Research. All data and materials will be available at github.com/ccnlab/deep-obj-cat upon publication.

Reprint requests should be sent to Randall C. O'Reilly, Department of Psychology, Computer Science, and Center for Neuroscience, University of California Davis, 1544 Newton Ct, Davis, CA 95618, or via e-mail: oreilly@ucdavis.edu.

Author Contributions

Randall C. O'Reilly: Conceptualization; Formal analysis; Funding acquisition; Investigation; Methodology; Project administration; Software; Supervision; Validation; Visualization; Writing – Original draft. Jacob L. Russin: Formal analysis; Investigation; Methodology; Validation; Writing – Review & editing. Maryam Zolfaghar: Investigation; Methodology; Validation; Writing – Review & editing. John Rohrlich: Conceptualization; Data curation; Investigation; Methodology; Software; Validation; Writing – Review & editing.

Funding Information

Randall C. O'Reilly: Office of Naval Research (<http://dx.doi.org/10.13039/1000000006>), grants D00014-12-C-0638, N00014-13-1-0067, N00014-14-1-0670, N00014-18-C-2067, and N00014-19-1-2684.

Diversity in Citation Practices

A retrospective analysis of the citations in every article published in this journal from 2010 to 2020 has revealed a persistent pattern of gender imbalance: Although the proportions of authorship teams (categorized by estimated

gender identification of first author/last author) publishing in the *Journal of Cognitive Neuroscience (JoCN)* during this period were $M(\text{an})/M = .408$, $W(\text{oman})/M = .335$, $M/W = .108$, and $W/W = .149$, the comparable proportions for the articles that these authorship teams cited were $M/M = .579$, $W/M = .243$, $M/W = .102$, and $W/W = .076$ (Fulvio et al., *JoCN*, 33:1, pp. 3–7). Consequently, *JoCN* encourages all authors to consider gender balance explicitly when selecting which articles to cite and gives them the opportunity to report their article's gender citation balance.

REFERENCES

- Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275, 220. DOI: <https://doi.org/10.1126/science.275.5297.221>, PMID: 8985017
- Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9, 147–169. DOI: https://doi.org/10.1207/s15516709cog0901_7
- Antonov, P. A., Chakravarthi, R., & Andersen, S. K. (2020). Too little, too late, and in the wrong place: Alpha band activity does not reflect an active mechanism of selective attention. *Neuroimage*, 219, 117006. DOI: <https://doi.org/10.1016/j.neuroimage.2020.117006>, PMID: 32485307
- Arcaro, M. J., Pinsk, M. A., & Kastner, S. (2015). The anatomical and functional organization of the human visual pulvinar. *Journal of Neuroscience*, 35, 9848–9871. DOI: <https://doi.org/10.1523/JNEUROSCI.1575-14.2015>
- Ashby, F. G., & Maddox, W. T. (2011). Human Category Learning 2.0. *Annals of the New York Academy of Sciences*, 1224, 147–161. DOI: <https://doi.org/10.1111/j.1749-6632.2010.05874.x>, PMID: 21182535, PMCID: PMC3076539
- Barczak, A., O'Connell, M. N., McGinnis, T., Ross, D., Mowery, T., Falchier, A., et al. (2018). Top-down, contextual entrainment of neuronal oscillations in the auditory thalamocortical circuit. *Proceedings of the National Academy of Sciences, U.S.A.*, 115, E7605–E7614. DOI: <https://doi.org/10.1073/pnas.1714684115>, PMID: 30037997, PMCID: PMC6094129
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76, 695–711. DOI: <https://doi.org/10.1016/j.neuron.2012.10.038>, PMID: 23177956, PMCID: PMC3777738
- Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., et al. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85, 390–401. DOI: <https://doi.org/10.1016/j.neuron.2014.12.018>, PMID: 25556836
- Bednar, J. A. (2012). Building a mechanistic model of the development and function of the primary visual cortex. *Journal of Physiology, Paris*, 106, 194–211. DOI: <https://doi.org/10.1016/j.jphysparis.2011.12.001>, PMID: 22343520
- Bender, D. B. (1982). Receptive-field properties of neurons in the macaque inferior pulvinar. *Journal of Neurophysiology*, 48, 1–17. DOI: <https://doi.org/10.1152/jn.1982.48.1.1>, PMID: 7119838
- Bender, D. B., & Youakim, M. (2001). Effect of attentive fixation in macaque thalamus and cortex. *Journal of Neurophysiology*, 85, 219–234. DOI: <https://doi.org/10.1152/jn.2001.85.1.219>, PMID: 11152722
- Bengio, Y., Mesnard, T., Fischer, A., Zhang, S., & Wu, Y. (2017). STDP-compatible approximation of backpropagation in an energy-based model. *Neural Computation*, 29, 555–577. DOI: https://doi.org/10.1162/NECO_a_00934, PMID: 28095200
- Bengio, Y., Yao, L., Alain, G., & Vincent, P. (2013). Generalized denoising auto-encoders as generative models. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems 26* (pp. 899–907). Curran Associates, Inc. <http://papers.nips.cc/paper/5023-generalized-denoising-auto-encoders-as-generative-models.pdf>
- Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 87, 527–570. DOI: <https://doi.org/10.1007/BF01797193>
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2, 32–48. DOI: <https://doi.org/10.1523/JNEUROSCI.02-01-00032.1982>, PMID: 7054394, PMCID: PMC6564292
- Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185–205). Cambridge, MA: MIT Press.
- Bortone, D. S., Olsen, S. R., & Scanziani, M. (2014). Translaminar inhibitory cells recruited by layer 6 corticothalamic neurons suppress visual cortex. *Neuron*, 82, 474–485. DOI: <https://doi.org/10.1016/j.neuron.2014.02.021>, PMID: 24656931, PMCID: PMC4068343
- Bourne, J. A., & Rosa, M. G. P. (2006). Hierarchical development of the primate visual cortex, as revealed by neurofilament immunoreactivity: Early maturation of the middle temporal area (MT). *Cerebral Cortex*, 16, 405–414. DOI: <https://doi.org/10.1093/cercor/bhi119>, PMID: 15944371
- Brette, R., & Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*, 94, 3637–3642. DOI: <https://doi.org/10.1152/jn.00686.2005>, PMID: 16014787
- Bridge, H., Leopold, D. A., & Bourne, J. A. (2016). Adaptive pulvinar circuitry supports visual cognition. *Trends in Cognitive Sciences*, 20, 146–157. DOI: <https://doi.org/10.1016/j.tics.2015.10.003>, PMID: 26553222, PMCID: PMC4724498
- Buffalo, E. A., Fries, P., Landman, R., Buschman, T. J., & Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *Proceedings of the National Academy of Sciences, U.S.A.*, 108, 11262–11267. DOI: <https://doi.org/10.1073/pnas.1011284108>, PMID: 21690410, PMCID: PMC3131344
- Busch, N. A., Dubois, J., & VanRullen, R. (2009). The phase of ongoing EEG oscillations predicts visual perception. *Journal of Neuroscience*, 29, 7869–7876. DOI: <https://doi.org/10.1523/JNEUROSCI.0113-09.2009>, PMID: 19535598, PMCID: PMC6665641
- Buzsáki, G. (2005). Theta rhythm of navigation: Link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, 15, 827–840. DOI: <https://doi.org/10.1002/hipo.20113>, PMID: 16149082
- Cadiou, C. F., Hong, H., Yamins, D. L. K., Pinto, N., Ardila, D., Solomon, E. A., et al. (2014). Deep neural networks rival the representation of primate IT cortex for core visual object recognition. *PLoS Computational Biology*, 10, e1003963. DOI: <https://doi.org/10.1371/journal.pcbi.1003963>, PMID: 25521294, PMCID: PMC4270441
- Cavanagh, P., Hunt, A. R., Afraz, A., & Rolfs, M. (2010). Visual stability based on remapping of attention pointers. *Trends in Cognitive Sciences*, 14, 147–153. DOI: <https://doi.org/10.1016/j.tics.2010.01.007>, PMID: 20189870, PMCID: PMC2847621
- Chaudhuri, R., Knoblauch, K., Gariel, M.-A., Kennedy, H., & Wang, X.-J. (2015). A large-scale circuit mechanism for

- hierarchical dynamical processing in the primate cortex. *Neuron*, 88, 419–431. **DOI:** <https://doi.org/10.1016/j.neuron.2015.09.008>, **PMID:** 26439530, **PMCID:** PMC4630024
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–204. **DOI:** <https://doi.org/10.1017/S0140525X12000477>, **PMID:** 23663408
- Clayton, M. S., Yeung, N., & Kadosh, R. C. (2018). The many characters of visual alpha oscillations. *European Journal of Neuroscience*, 48, 2498–2508. **DOI:** <https://doi.org/10.1111/ejn.13747>, **PMID:** 29044823
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, 120, 235–253. **DOI:** <https://doi.org/10.1037/0096-3445.120.3.235>, **PMID:** 1836490
- Colby, C. L., Duhamel, J. R., & Goldberg, M. E. (1997). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *Journal of Neurophysiology*, 76, 2841–2852. **DOI:** <https://doi.org/10.1152/jn.1996.76.5.2841>, **PMID:** 8930237
- Connors, B. W., Gutnick, M. J., & Prince, D. A. (1982). Electrophysiological properties of neocortical neurons in vitro. *Journal of Neurophysiology*, 48, 1302–1320. **DOI:** <https://doi.org/10.1152/jn.1982.48.6.1302>, **PMID:** 6296328
- Cooper, L. N., & Bear, M. F. (2012). The BCM theory of synapse modification at 30: Interaction of theory with experiment. *Nature Reviews Neuroscience*, 13, 798–810. **DOI:** <https://doi.org/10.1038/nrn3353>, **PMID:** 23080416
- Crick, F. (1984). Function of the thalamic reticular complex: The searchlight hypothesis. *Proceedings of the National Academy of Sciences, USA*, 81, 4586–4590. **DOI:** <https://doi.org/10.1073/pnas.81.14.4586>, **PMID:** 6589612, **PMCID:** PMC345636
- Crick, F. (1989). The recent excitement about neural networks. *Nature*, 337, 129–132. **DOI:** <https://doi.org/10.1038/337129a0>, **PMID:** 2911347
- Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5, 613–624. **DOI:** <https://doi.org/10.1162/neco.1993.5.4.613>
- Dayan, P., Hinton, G. E., Neal, R. N., & Zemel, R. S. (1995). The Helmholtz machine. *Neural Computation*, 7, 889–904. **DOI:** <https://doi.org/10.1162/neco.1995.7.5.889>, **PMID:** 7584891
- de Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, 22, 764–779. **DOI:** <https://doi.org/10.1016/j.tics.2018.06.002>, **PMID:** 30122170
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222. **DOI:** <https://doi.org/10.1146/annurev.ne.18.030195.001205>, **PMID:** 7605061
- Duhamel, J. R., Colby, C. L., & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255, 90–92. **DOI:** <https://doi.org/10.1126/science.1553535>, **PMID:** 1553535
- Elman, J., Bates, E., Karmiloff-Smith, A., Johnson, M., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press. **DOI:** <https://doi.org/10.7551/mitpress/5929.001.0001>
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211. **DOI:** [https://doi.org/10.1016/0364-0213\(90\)90002-E](https://doi.org/10.1016/0364-0213(90)90002-E)
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47. **DOI:** <https://doi.org/10.1093/cercor/1.1.1>, **PMID:** 1822724
- Fiebelkorn, I. C., & Kastner, S. (2019). A rhythmic theory of attention. *Trends in Cognitive Sciences*, 23, 87–101. **DOI:** <https://doi.org/10.1016/j.tics.2018.11.009>, **PMID:** 30591373, **PMCID:** PMC6343831
- Fiebelkorn, I. C., Pinski, M. A., & Kastner, S. (2018). A dynamic interplay within the frontoparietal network underlies rhythmic spatial attention. *Neuron*, 99, 842–853. **DOI:** <https://doi.org/10.1016/j.neuron.2018.07.038>, **PMID:** 30138590, **PMCID:** PMC6474777
- Fiser, A., Mahringer, D., Oyibo, H. K., Petersen, A. V., Leinweber, M., & Keller, G. B. (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience*, 19, 1658–1664. **DOI:** <https://doi.org/10.1038/nn.4385>, **PMID:** 27618309
- Foldiak, P. (1991). Learning invariance from transformation sequences. *Neural Computation*, 3, 194–200. **DOI:** <https://doi.org/10.1162/neco.1991.3.2.194>, **PMID:** 31167302
- Foster, J. J., & Awh, E. (2019). The role of alpha oscillations in spatial attention: Limited evidence for a suppression account. *Current Opinion in Psychology*, 29, 34–40. **DOI:** <https://doi.org/10.1016/j.copsyc.2018.11.001>, **PMID:** 30472541, **PMCID:** PMC6506396
- Franceschetti, S., Guatteo, E., Panzica, F., Sancini, G., Wanke, E., & Avanzini, G. (1995). Ionic mechanisms underlying burst firing in pyramidal neurons: Intracellular study in rat sensorimotor cortex. *Brain Research*, 696, 127–139. **DOI:** [https://doi.org/10.1016/0006-8993\(95\)00807-3](https://doi.org/10.1016/0006-8993(95)00807-3), **PMID:** 8574660
- Fries, P., Womelsdorf, T., Oostenveld, R., & Desimone, R. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *Journal of Neuroscience*, 28, 4823–4835. **DOI:** <https://doi.org/10.1523/JNEUROSCI.4499-07.2008>, **PMID:** 18448659, **PMCID:** PMC3844818
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 360, 815–836. **DOI:** <https://doi.org/10.1098/rstb.2005.1622>, **PMID:** 15937014, **PMCID:** PMC1569488
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138. **DOI:** <https://doi.org/10.1038/nrn2787>, **PMID:** 20068583
- Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, 37, 66–74. **DOI:** <https://doi.org/10.1016/j.conb.2016.01.010>, **PMID:** 26851755
- Gardner, M. P. H., Schoenbaum, G., & Gershman, S. J. (2018). Rethinking dopamine as generalized prediction error. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, 285, 20181645. **DOI:** <https://doi.org/10.1098/rspb.2018.1645>, **PMID:** 30464063, **PMCID:** PMC6253385
- Gavornik, J. P., & Bear, M. F. (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nature Neuroscience*, 17, 732–737. **DOI:** <https://doi.org/10.1038/nn.3683>, **PMID:** 24657967, **PMCID:** PMC4167369
- George, D., & Hawkins, J. (2009). Towards a mathematical theory of cortical micro-circuits. *PLoS Computational Biology*, 5, e1000532. **DOI:** <https://doi.org/10.1371/journal.pcbi.1000532>, **PMID:** 19816557, **PMCID:** PMC2749218
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15, 20–25. **DOI:** [https://doi.org/10.1016/0166-2236\(92\)90344-8](https://doi.org/10.1016/0166-2236(92)90344-8), **PMID:** 1374953
- Gottlieb, J. P., Kusunoki, M., & Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, 391, 481–484. **DOI:** <https://doi.org/10.1038/35135>, **PMID:** 9461214
- Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects.

- Trends in Cognitive Sciences*, 10, 14–23. DOI: <https://doi.org/10.1016/j.tics.2005.11.006>, PMID: 16321563
- Grossberg, S. (1999). How does the cerebral cortex work? Learning, attention, and grouping by the laminar circuits of visual cortex. *Spatial Vision*, 12, 163–185. DOI: <https://doi.org/10.1163/156856899X00102>, PMID: 10221426
- Gruber, W. R., Klimesch, W., Sauseng, P., & Doppelmayr, M. (2005). Alpha phase synchronization predicts P1 and N1 latency and amplitude size. *Cerebral Cortex*, 15, 371–377. DOI: <https://doi.org/10.1093/cercor/bhh139>, PMID: 15749980
- Gundlach, C., Moratti, S., Forschack, N., & Müller, M. M. (2020). Spatial attentional selection modulates early visual stimulus processing independently of visual alpha modulations. *Cerebral Cortex*, 30, 3686–3703. DOI: <https://doi.org/10.1093/cercor/bhz335>, PMID: 31907512
- Halassa, M. M., & Kastner, S. (2017). Thalamic functions in distributed cognitive control. *Nature Neuroscience*, 20, 1669. DOI: <https://doi.org/10.1038/s41593-017-0020-1>, PMID: 29184210
- Harris, K. D., & Shepherd, G. M. G. (2015). The neocortical circuit: Themes and variations. *Nature Neuroscience*, 18, 170–181. DOI: <https://doi.org/10.1038/nn.3917>, PMID: 25622573, PMCID: PMC4889215
- Hawkins, J., & Blakeslee, S. (2004). *On intelligence*. New York: Times Books.
- Hennig, M. H. (2013). Theoretical models of synaptic short term plasticity. *Frontiers in Computational Neuroscience*, 7, 45. DOI: <https://doi.org/10.3389/fncom.2013.00154>, PMID: 24198783, PMCID: PMC3812535
- Hinton, G. E., & McClelland, J. L. (1988). Learning representations by recirculation. In D. Z. Anderson (Ed.), *Neural information processing systems (NIPS 1987)* (pp. 358–366). New York: American Institute of Physics. <http://papers.nips.cc/paper/78-learning-representations-by-recirculation.pdf>
- Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313, 504–507. DOI: <https://doi.org/10.1126/science.1127647>, PMID: 16873662
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109, 679–709. DOI: <https://doi.org/10.1037/0033-295X.109.4.679>, PMID: 12374324
- Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, U.S.A.*, 81, 3088–3092. DOI: <https://doi.org/10.1073/pnas.81.10.3088>, PMID: 6587342, PMCID: PMC345226
- Issa, E. B., Cadieu, C. F., & DiCarlo, J. J. (2018). Neural dynamics at successive stages of the ventral visual stream are consistent with hierarchical error signals. *eLife*, 7, e42870. DOI: <https://doi.org/10.7554/eLife.42870>, PMID: 30484773, PMCID: PMC6296785
- Jaegle, A., & Ro, T. (2013). Direct control of visual perception with phase-specific modulation of posterior parietal cortex. *Journal of Cognitive Neuroscience*, 26, 422–432. DOI: https://doi.org/10.1162/jocn_a_00494, PMID: 24116843
- Jaramillo, J., Mejias, J. F., & Wang, X.-J. (2019). Engagement of pulvino-cortical feedforward and feedback pathways in cognitive computations. *Neuron*, 101, 321–336. DOI: <https://doi.org/10.1016/j.neuron.2018.11.023>, PMID: 30553546, PMCID: PMC6650151
- Jensen, O., Bonnefond, M., Marshall, T. R., & Tiesinga, P. (2015). Oscillatory mechanisms of feedforward and feedback visual processing. *Trends in Neurosciences*, 38, 192–194. DOI: <https://doi.org/10.1016/j.tics.2015.02.006>, PMID: 25765320
- Jensen, O., Bonnefond, M., & VanRullen, R. (2012). An oscillatory mechanism for prioritizing salient unattended stimuli. *Trends in Cognitive Sciences*, 16, 200–206. DOI: <https://doi.org/10.1016/j.tics.2012.03.002>, PMID: 22436764
- Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: Gating by inhibition. *Frontiers in Human Neuroscience*, 4, 186. DOI: <https://doi.org/10.3389/fnhum.2010.00186>, PMID: 21119777, PMCID: PMC2990626
- Jordan, M. I. (1989). Serial order: A parallel, distributed processing approach. In J. L. Elman & D. E. Rumelhart (Eds.), *Advances in connectionist theory: Speech*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kachergis, G., Wyatte, D., O'Reilly, R. C., de Kleijn, R., & Hommel, B. (2014). A continuous-time neural model for sequential action. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 369, 20130623. DOI: <https://doi.org/10.1098/rstb.2013.0623>, PMID: 25267830, PMCID: PMC4186241
- Kahana, M. J., Seelig, D., & Madsen, J. R. (2001). Theta returns. *Current Opinion in Neurobiology*, 11, 739–744. DOI: [https://doi.org/10.1016/S0959-4388\(01\)00278-1](https://doi.org/10.1016/S0959-4388(01)00278-1), PMID: 11741027
- Kawato, M., Hayakawa, H., & Inui, T. (1993). A forward-inverse optics model of reciprocal connections between visual cortical areas. *Network: Computation in Neural Systems*, 4, 415–422. DOI: https://doi.org/10.1088/0954-898X_4_4_001
- Keitel, C., Keitel, A., Benwell, C. S. Y., Daube, C., Thut, G., & Gross, J. (2019). Stimulus-driven brain rhythms within the alpha band: The attentional-modulation conundrum. *Journal of Neuroscience*, 39, 3119–3129. DOI: <https://doi.org/10.1523/JNEUROSCI.1633-18.2019>, PMID: 30770401, PMCID: PMC6468105
- Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. J. (2006). Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained visuospatial attention. *Journal of Neurophysiology*, 95, 3844–3851. DOI: <https://doi.org/10.1152/jn.01234.2005>, PMID: 16571739
- Khaligh-Razavi, S.-M., & Kriegeskorte, N. (2014). Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Computational Biology*, 10, e1003915. DOI: <https://doi.org/10.1371/journal.pcbi.1003915>, PMID: 25375136, PMCID: PMC4222664
- Kiorpes, L., Price, T., Hall-Haro, C., & Movshon, J. A. (2012). Development of sensitivity to global form and motion in macaque monkeys (*Macaca nemestrina*). *Vision Research*, 63, 34–42. DOI: <https://doi.org/10.1016/j.visres.2012.04.018>, PMID: 22580018, PMCID: PMC3374036
- Klimesch, W. (2011). Evoked alpha and early access to the knowledge system: The P1 inhibition timing hypothesis. *Brain Research*, 1408, 52–71. DOI: <https://doi.org/10.1016/j.brainres.2011.06.003>, PMID: 21774917, PMCID: PMC3158852
- Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: The inhibition-timing hypothesis. *Brain Research Reviews*, 53, 63–88. DOI: <https://doi.org/10.1016/j.brainresrev.2006.06.003>, PMID: 16887192
- Kobatake, E., & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway. *Journal of Neurophysiology*, 71, 856–867. DOI: <https://doi.org/10.1152/jn.1994.71.3.856>, PMID: 8201425
- Kogo, N., & Trengove, C. (2015). Is predictive coding theory articulated enough to be testable? *Frontiers in Computational Neuroscience*, 9, 111. DOI: <https://doi.org/10.3389/fncom.2015.00111>, PMID: 26441621, PMCID: PMC4561670
- Kok, P., & de Lange, F. P. (2015). Predictive coding in sensory cortex. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience*

- (pp. 221–244). New York: Springer. DOI: <https://doi.org/10.1007/978-1-4939-2236-9>
- Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less is more: Expectation sharpens representations in the primary visual cortex. *Neuron*, *75*, 265–270. DOI: <https://doi.org/10.1016/j.neuron.2012.04.034>, PMID: 22841311
- Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., & Miyamoto, A. (2013). Responses of pulvinar neurons reflect a subject's confidence in visual categorization. *Nature Neuroscience*, *16*, 749–755. DOI: <https://doi.org/10.1038/nn.3393>, PMID: 23666179
- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis—Connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*, 4. DOI: <https://doi.org/10.3389/fnro.06.004.2008>, PMID: 19104670, PMCID: PMC2605405
- LaBerge, D., & Buchsbaum, M. S. (1990). Positron emission tomographic measurements of pulvinar activity during an attention task. *Journal of Neuroscience*, *10*, 613–619. DOI: <https://doi.org/10.1523/JNEUROSCI.10-02-00613.1990>, PMID: 2303863, PMCID: PMC6570168
- Larkum, M. E., Zhu, J. J., & Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature*, *398*, 338–341. DOI: <https://doi.org/10.1038/18686>, PMID: 10192334
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*, 436–444. DOI: <https://doi.org/10.1038/nature14539>, PMID: 26017442
- Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America*, *20*, 1434–1448. DOI: <https://doi.org/10.1364/JOSAA.20.001434>, PMID: 12868647
- Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*, *21*, 335–346. DOI: <https://doi.org/10.1038/s41583-020-0277-3>, PMID: 32303713
- Lim, S., McKee, J. L., Woloszyn, L., Amit, Y., Freedman, D. J., Sheinberg, D. L., et al. (2015). Inferring learning rules from distributions of firing rates in cortical neurons. *Nature Neuroscience*, *18*, 1804–1810. DOI: <https://doi.org/10.1038/nn.4158>, PMID: 26523643, PMCID: PMC4666720
- Lotter, W., Kreiman, G., & Cox, D. (2016). Deep predictive coding networks for video prediction and unsupervised learning. *arXiv:1605.08104 [cs, q-bio]*. <http://arxiv.org/abs/1605.08104>
- Luczak, A., Bartho, P., & Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, *62*, 413–425. DOI: <https://doi.org/10.1016/j.neuron.2009.03.014>, PMID: 19447096, PMCID: PMC2696272
- Luczak, A., Bartho, P., & Harris, K. D. (2013). Gating of sensory input by spontaneous cortical activity. *Journal of Neuroscience*, *33*, 1684–1695. DOI: <https://doi.org/10.1523/JNEUROSCI.2928-12.2013>, PMID: 23345241, PMCID: PMC3672963
- Lüscher, C., & Malenka, R. C. (2012). NMDA receptor-dependent long-term potentiation and long-term depression (LTP/LTD). *Cold Spring Harbor Perspectives in Biology*, *4*, a005710. DOI: <https://doi.org/10.1101/cshperspect.a005710>, PMID: 22510460, PMCID: PMC3367554
- Maier, A., Adams, G. K., Aura, C., & Leopold, D. A. (2010). Distinct superficial and deep laminar domains of activity in the visual cortex during rest and stimulation. *Frontiers in Systems Neuroscience*, *4*, 31. DOI: <https://doi.org/10.3389/fnro.2010.00031>, PMID: 20802856, PMCID: PMC2928665
- Maier, A., Aura, C. J., & Leopold, D. A. (2011). Infragranular sources of sustained local field potential responses in macaque primary visual cortex. *Journal of Neuroscience*, *31*, 1971–1980. DOI: <https://doi.org/10.1523/JNEUROSCI.5300-09.2011>, PMID: 21307235, PMCID: PMC3075009
- Makeig, S., Westerfield, M., Jung, T. P., Enghoff, S., Townsend, J., Courchesne, E., et al. (2002). Dynamic brain sources of visual evoked responses. *Science*, *295*, 690–693. DOI: <https://doi.org/10.1126/science.1066168>, PMID: 11809976
- Marino, A. C., & Mazer, J. A. (2016). Perisaccadic updating of visual representations and attentional states: Linking behavior and neurophysiology. *Frontiers in Systems Neuroscience*, *10*, 3. DOI: <https://doi.org/10.3389/fnro.2016.00003>, PMID: 26903820, PMCID: PMC4743436
- Markov, N. T., Ercsey-Ravasz, M. M., Ribeiro Gomes, A. R., Lamy, C., Magrou, L., Vezoli, J., et al. (2014). A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cerebral Cortex*, *24*, 17–36. DOI: <https://doi.org/10.1093/cercor/bhs270>, PMID: 23010748, PMCID: PMC3862262
- Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., et al. (2014). Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex: Cortical counterstreams. *Journal of Comparative Neurology*, *522*, 225–259. DOI: <https://doi.org/10.1002/cne.23458>, PMID: 23983048, PMCID: PMC4255240
- Martinez-Conde, S., Macknik, S. L., & Hubel, D. H. (2004). The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience*, *5*, 229–240. DOI: <https://doi.org/10.1038/nrn1348>, PMID: 14976522
- Martinez-Conde, S., Otero-Millan, J., & Macknik, S. L. (2013). The impact of microsaccades on vision: Towards a unified theory of saccadic function. *Nature Reviews Neuroscience*, *14*, 83–96. DOI: <https://doi.org/10.1038/nrn3405>, PMID: 23329159
- Mathewson, K., Gratton, G., Fabiani, M., Beck, D., & Ro, T. (2009). To see or not to see: Prestimulus alpha phase predicts visual awareness. *Journal of Neuroscience*, *29*, 2725–2732. DOI: <https://doi.org/10.1523/JNEUROSCI.3963-08.2009>, PMID: 19261866, PMCID: PMC2724892
- Mathewson, K. E., Fabiani, M., Gratton, G., Beck, D. M., & Lleras, A. (2010). Rescuing stimuli from invisibility: Inducing a momentary release from visual masking with pre-target entrainment. *Cognition*, *115*, 186–191. DOI: <https://doi.org/10.1016/j.cognition.2009.11.010>, PMID: 20035933
- Mathewson, K. E., Prudhomme, C., Fabiani, M., Beck, D. M., Lleras, A., & Gratton, G. (2012). Making waves in the stream of consciousness: Entraining oscillations in EEG alpha and fluctuations in visual awareness with rhythmic visual stimulation. *Journal of Cognitive Neuroscience*, *24*, 2321–2333. DOI: https://doi.org/10.1162/jocn_a_00288, PMID: 22905825
- Mayer, A., Schwiedrzik, C. M., Wibrall, M., Singer, W., & Melloni, L. (2016). Expecting to see a letter: Alpha oscillations as carriers of top-down sensory predictions. *Cerebral Cortex*, *26*, 3146–3160. DOI: <https://doi.org/10.1093/cercor/bhv146>, PMID: 26142463
- Meyer, T., & Olson, C. R. (2011). Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *108*, 19401–19406. DOI: <https://doi.org/10.1073/pnas.1112895108>, PMID: 22084090, PMCID: PMC3228439
- Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J.-M., Kennedy, H., & Fries, P. (2016). Alpha-beta and gamma rhythms subserve feedback and feedforward influences among human visual cortical areas. *Neuron*, *89*, 384–397. DOI: <https://doi.org/10.1016/j.neuron.2015.12.018>, PMID: 26777277, PMCID: PMC4871751
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202. DOI: <https://doi.org/10.1146/annurev.neuro.24.1.167>, PMID: 11283309
- Miller, K. D. (1994). A model for the development of simple cell receptive fields and the ordered arrangement of orientation

- columns through activity-dependent competition between ON- and OFF-center inputs. *Journal of Neuroscience*, *14*, 409–441. **DOI:** <https://doi.org/10.1523/JNEUROSCI.14-01-00409.1994>, **PMID:** 8283248, **PMCID:** PMC6576834
- Müller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, *285*, 1405–1408. **DOI:** <https://doi.org/10.1126/science.285.5432.1405>, **PMID:** 10464100
- Mumford, D. (1991). On the computational architecture of the neocortex. *Biological Cybernetics*, *65*, 135–145. **DOI:** <https://doi.org/10.1007/BF00202389>, **PMID:** 1912004
- Mumford, D. (1992). On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biological Cybernetics*, *66*, 241–251. **DOI:** <https://doi.org/10.1007/BF00198477>, **PMID:** 1540675
- Nakamura, K., & Colby, C. L. (2002). Updating of the visual representation in monkey striate and extrastriate cortex during saccades. *Proceedings of the National Academy of Sciences, U.S.A.*, *99*, 4026–4031. **DOI:** <https://doi.org/10.1073/pnas.052379899>, **PMID:** 11904446, **PMCID:** PMC122642
- Neupane, S., Guitton, D., & Pack, C. C. (2016). Two distinct types of remapping in primate cortical area V4. *Nature Communications*, *7*, 10402. **DOI:** <https://doi.org/10.1038/ncomms10402>, **PMID:** 26832423, **PMCID:** PMC4740356
- Neupane, S., Guitton, D., & Pack, C. C. (2017). Coherent alpha oscillations link current and future receptive fields during saccades. *Proceedings of the National Academy of Sciences, U.S.A.*, *114*, E5979–E5985. **DOI:** <https://doi.org/10.1073/pnas.1701672114>, **PMID:** 28673993, **PMCID:** PMC5530666
- Neupane, S., Guitton, D., & Pack, C. C. (2020). Perisaccadic remapping: What? how? why? *Reviews in the Neurosciences*, *31*, 505–520. **DOI:** <https://doi.org/10.1515/revneuro-2019-0097>, **PMID:** 32242834
- Nunn, C. M. H., & Osselson, J. W. (1974). The influence of the EEG alpha rhythm on the perception of visual stimuli. *Psychophysiology*, *11*, 294–303. **DOI:** <https://doi.org/10.1111/j.1469-8986.1974.tb00547.x>, **PMID:** 4421317
- O’Herron, P., & von der Heydt, R. (2013). Remapping of border ownership in the visual cortex. *Journal of Neuroscience*, *33*, 1964–1974. **DOI:** <https://doi.org/10.1523/JNEUROSCI.2797-12.2013>, **PMID:** 23365235, **PMCID:** PMC4086328
- Olsen, S., Bortone, D., Adesnik, H., & Scanziani, M. (2012). Gain control by layer six in cortical circuits of vision. *Nature*, *483*, 47–52. **DOI:** <https://doi.org/10.1038/nature10835>, **PMID:** 22367547, **PMCID:** PMC3636977
- O’Reilly, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, *8*, 895–938. **DOI:** <https://doi.org/10.1162/neco.1996.8.5.895>
- O’Reilly, R. C. (1998). Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences*, *2*, 455–462. **DOI:** [https://doi.org/10.1016/S1364-6613\(98\)01241-8](https://doi.org/10.1016/S1364-6613(98)01241-8)
- O’Reilly, R. C., Hazy, T. E., & Herd, S. A. (2016). The Leabra cognitive architecture: How to play 20 principles with nature and win! In S. Chipman (Ed.), *Oxford handbook of cognitive science*. Oxford, UK: Oxford University Press. <http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199842193.001.0001/oxfordhb-9780199842193-e-8>. **DOI:** <https://doi.org/10.1093/oxfordhb/9780199842193.013.8>
- O’Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press. **DOI:** <https://doi.org/10.7551/mitpress/2014.001.0001>
- O’Reilly, R. C., Munakata, Y., Frank, M. J., Hazy, T. E., & Contributors. (2012). *Computational cognitive neuroscience* (1st ed.). Wiki Book. <http://ccnbook.colorado.edu>
- O’Reilly, R. C., Wyatte, D., Herd, S., Mingus, B., & Jilk, D. J. (2013). Recurrent processing during object recognition. *Frontiers in Psychology*, *4*, 124. **DOI:** <https://doi.org/10.3389/fpsyg.2013.00124>, **PMID:** 23554596, **PMCID:** PMC3612699
- O’Reilly, R. C., Wyatte, D., & Rohrlich, J. (2014). Learning through time in the thalamocortical loops. *arXiv:1407.3432 [q-bio]*. <http://arxiv.org/abs/1407.3432>
- O’Reilly, R. C., Wyatte, D. R., & Rohrlich, J. (2017). Deep predictive learning: A comprehensive model of three visual streams. *arXiv:1709.04654 [q-bio]*. <http://arxiv.org/abs/1709.04654>
- Ouden, H. E. M., Kok, P., & Lange, F. P. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology*, *3*, 548. **DOI:** <https://doi.org/10.3389/fpsyg.2012.00548>, **PMID:** 23248610, **PMCID:** PMC3518876
- Palva, S., & Palva, J. M. (2011). Functional roles of alpha-band phase synchronization in local and large-scale cortical networks. *Frontiers in Psychology*, *2*, 204. **DOI:** <https://doi.org/10.3389/fpsyg.2011.00204>, **PMID:** 21922012, **PMCID:** PMC3166799
- Pennartz, C. M., Dora, S., Muckli, L., & Lorteije, J. A. (2019). Towards a unified view on pathways and functions of neural recurrent processing. *Trends in Neurosciences*, *42*, 589–603. **DOI:** <https://doi.org/10.1016/j.tins.2019.07.005>, **PMID:** 31399289
- Petersen, S. E., Robinson, D. L., & Keys, W. (1985). Pulvinar nuclei of the behaving rhesus monkey: Visual responses and their modulation. *Journal of Neurophysiology*, *54*, 867–886. **DOI:** <https://doi.org/10.1152/jn.1985.54.4.867>, **PMID:** 4067625
- Petrof, I., Viaene, A. N., & Sherman, S. M. (2012). Two populations of corticothalamic and interareal corticocortical cells in the subgranular layers of the mouse primary sensory cortices. *Journal of Comparative Neurology*, *520*, 1678–1686. **DOI:** <https://doi.org/10.1002/cne.23006>, **PMID:** 22120996, **PMCID:** PMC3561675
- Pinault, D. (2004). The thalamic reticular nucleus: Structure, function and concept. *Brain Research*, *46*, 1–31. **DOI:** <https://doi.org/10.1016/j.brainresrev.2004.04.008>, **PMID:** 15297152
- Pineda, F. J. (1987). Generalization of backpropagation to recurrent neural networks. *Physical Review Letters*, *18*, 2229–2232. **DOI:** <https://doi.org/10.1103/PhysRevLett.59.2229>, **PMID:** 10035458
- Pouget, A., & Sejnowski, T. J. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience*, *9*, 222–237. **DOI:** <https://doi.org/10.1162/jocn.1997.9.2.222>, **PMID:** 23962013
- Privman, E., Malach, R., & Yeshurun, Y. (2013). Modeling the electrical field created by mass neural activity. *Neural Networks*, *40*, 44–51. **DOI:** <https://doi.org/10.1016/j.neunet.2013.01.004>, **PMID:** 23391515
- Purushothaman, G., Marion, R., Li, K., & Casagrande, V. A. (2012). Gating and control of primary visual cortex by pulvinar. *Nature Neuroscience*, *15*, 905–912. **DOI:** <https://doi.org/10.1038/nn.3106>, **PMID:** 22561455, **PMCID:** PMC3430824
- Pylshyn, Z. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition*, *32*, 65–97. **DOI:** [https://doi.org/10.1016/0010-0277\(89\)90014-0](https://doi.org/10.1016/0010-0277(89)90014-0)
- Rajalingham, R., Issa, E. B., Bashivan, P., Kar, K., Schmidt, K., & DiCarlo, J. J. (2018). Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *Journal of Neuroscience*, *38*, 7255–7269. **DOI:** <https://doi.org/10.1101/240614>, **PMID:** 30006365, **PMCID:** PMC6096043
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*, 79–87. **DOI:** <https://doi.org/10.1038/4580>, **PMID:** 10195184

- Ray, S., & Maunsell, J. H. R. (2011). Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biology*, *9*, e1000610. **DOI:** <https://doi.org/10.1371/journal.pbio.1000610>, **PMID:** 21532743, **PMCID:** PMC3075230
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, *6*, 855–863. **DOI:** [https://doi.org/10.1016/S0022-5371\(67\)80149-X](https://doi.org/10.1016/S0022-5371(67)80149-X)
- Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, *19*, 1736–1753. **DOI:** <https://doi.org/10.1523/JNEUROSCI.19-05-01736.1999>, **PMID:** 10024360, **PMCID:** PMC6782185
- Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61*, 168–185. **DOI:** <https://doi.org/10.1016/j.neuron.2009.01.002>, **PMID:** 19186161, **PMCID:** PMC2752446
- Richter, D., & de Lange, F. P. (2019). Statistical learning attenuates visual activity only for attended stimuli. *eLife*, *8*, e47869. **DOI:** <https://doi.org/10.7554/eLife.47869>, **PMID:** 31442202, **PMCID:** PMC6731093
- Robinson, D. L. (1993). Functional contributions of the primate pulvinar. *Progress in Brain Research*, *95*, 371–380. **DOI:** [https://doi.org/10.1016/s0079-6123\(08\)60382-9](https://doi.org/10.1016/s0079-6123(08)60382-9), **PMID:** 8493346
- Rockland, K. S. (1996). Two types of corticopulvinar terminations: Round (type 2) and elongate (type 1). *Journal of Comparative Neurology*, *368*, 57–87. **DOI:** [https://doi.org/10.1002/\(sici\)1096-9861\(19960422\)368:1%3C57::aid-cne5%3E3.0.co;2-j](https://doi.org/10.1002/(sici)1096-9861(19960422)368:1%3C57::aid-cne5%3E3.0.co;2-j), **PMID:** 8725294
- Rockland, K. S. (1998). Convergence and branching patterns of round, type 2 corticopulvinar axons. *Journal of Comparative Neurology*, *390*, 515–536. **DOI:** [https://doi.org/10.1002/\(sici\)1096-9861\(19980126\)390:4%3C515::aid-cne5%3E3.0.co;2-3](https://doi.org/10.1002/(sici)1096-9861(19980126)390:4%3C515::aid-cne5%3E3.0.co;2-3), **PMID:** 9450533
- Rockland, K. S., & Pandya, D. N. (1979). Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Research*, *179*, 3–20. **DOI:** [https://doi.org/10.1016/0006-8993\(79\)90485-2](https://doi.org/10.1016/0006-8993(79)90485-2), **PMID:** 116716
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*, 533–536. **DOI:** <https://doi.org/10.1038/323533a0>
- Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: Part 2. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, *89*, 60–94. **DOI:** <https://doi.org/10.1037/0033-295X.89.1.60>, **PMID:** 7058229
- Saalmann, Y. B., & Kastner, S. (2011). Cognitive and perceptual functions of the visual thalamus. *Neuron*, *71*, 209–223. **DOI:** <https://doi.org/10.1016/j.neuron.2011.06.027>, **PMID:** 21791281, **PMCID:** PMC3148184
- Saalmann, Y. B., Pinsk, M. A., Wang, L., Li, X., & Kastner, S. (2012). The pulvinar regulates information transmission between cortical areas based on attention demands. *Science*, *337*, 753–756. **DOI:** <https://doi.org/10.1126/science.1223082>, **PMID:** 22879517, **PMCID:** PMC3714098
- Sakata, S., & Harris, K. D. (2009). Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron*, *64*, 404–418. **DOI:** <https://doi.org/10.1016/j.neuron.2009.09.020>, **PMID:** 19914188, **PMCID:** PMC2778614
- Sakata, S., & Harris, K. D. (2012). Laminar-dependent effects of cortical state on auditory cortical spontaneous activity. *Frontiers in Neural Circuits*, *6*, 109. **DOI:** <https://doi.org/10.3389/fncir.2012.00109>, **PMID:** 23267317, **PMCID:** PMC3527822
- Samaha, J., Bauer, P., Cimaroli, S., & Postle, B. R. (2015). Top-down control of the phase of alpha-band oscillations as a mechanism for temporal prediction. *Proceedings of the National Academy of Sciences, U.S.A.*, *112*, 8439–8444. **DOI:** <https://doi.org/10.1073/pnas.1503686112>, **PMID:** 26100913, **PMCID:** PMC4500260
- Sherman, M. T., Kanai, R., Seth, A. K., & VanRullen, R. (2016). Rhythmic influence of top-down perceptual priors in the phase of prestimulus occipital alpha oscillations. *Journal of Cognitive Neuroscience*, *28*, 1318–1330. **DOI:** https://doi.org/10.1162/jocn_a_00973, **PMID:** 27082046
- Sherman, S. M. (2014). The function of metabotropic glutamate receptors in thalamus and cortex. *Neuroscientist*, *20*, 146–149. **DOI:** <https://doi.org/10.1177/1073858413478490>, **PMID:** 23459618, **PMCID:** PMC4747429
- Sherman, S. M., & Guillery, R. W. (2006). *Exploring the thalamus and its role in cortical function*. Cambridge, MA: MIT Press. <http://www.scholarpedia.org/article/Thalamus>
- Sherman, S. M., & Guillery, R. W. (2011). Distinct functions for direct and transthalamic corticocortical connections. *Journal of Neurophysiology*, *106*, 1068–1077. **DOI:** <https://doi.org/10.1152/jn.00429.2011>, **PMID:** 21676936
- Sherman, S. M., & Guillery, R. W. (2013). *Functional connections of cortical areas: A new view from the thalamus*. Cambridge, MA: MIT Press. **DOI:** <https://doi.org/10.7551/mitpress/9780262019309.001.0001>
- Shipp, S. (2003). The functional logic of cortico-pulvinar connections. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *358*, 1605–1624. **DOI:** <https://doi.org/10.1098/rstb.2002.1213>, **PMID:** 14561322, **PMCID:** PMC1693262
- Shouval, H. Z. S., Bear, M. F., & Cooper, L. N. (2002). A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proceedings of the National Academy of Sciences, U.S.A.*, *99*, 10831–10836. **DOI:** <https://doi.org/10.1073/pnas.152343099>, **PMID:** 12136127, **PMCID:** PMC125058
- Shrager, J., & Johnson, M. H. (1996). Dynamic plasticity influences the emergence of function in a simple cortical array. *Neural Networks*, *9*, 1119–1129. **DOI:** [https://doi.org/10.1016/0893-6080\(96\)00033-0](https://doi.org/10.1016/0893-6080(96)00033-0), **PMID:** 12662587
- Silva, L. R., Amitai, Y., & Connors, B. W. (1991). Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Science*, *251*, 432–435. **DOI:** <https://doi.org/10.1126/science.1824881>, **PMID:** 1824881
- Snow, J. C., Allen, H. A., Rafal, R. D., & Humphreys, G. W. (2009). Impaired attentional selection following lesions to human pulvinar: Evidence for homology between human and monkey. *Proceedings of the National Academy of Sciences, U.S.A.*, *106*, 4054–4059. **DOI:** <https://doi.org/10.1073/pnas.0810086106>, **PMID:** 19237580, **PMCID:** PMC2656203
- Solis-Vivanco, R., Jensen, O., & Bonnefond, M. (2018). Top-down control of alpha phase adjustment in anticipation of temporally predictable visual stimuli. *Journal of Cognitive Neuroscience*, *30*, 1157–1169. **DOI:** https://doi.org/10.1162/jocn_a_01280, **PMID:** 29762100
- Solomon, E. A., Kragel, J. E., Sperling, M. R., Sharan, A., Worrell, G., Kucewicz, M., et al. (2017). Widespread theta synchrony and high-frequency desynchronization underlies enhanced cognition. *Nature Communications*, *8*, 1704. **DOI:** <https://doi.org/10.1038/s41467-017-01763-2>, **PMID:** 29167419, **PMCID:** PMC5700170
- Spaak, E., Bonnefond, M., Maier, A., Leopold, D. A., & Jensen, O. (2012). Layer-specific entrainment of gamma-band neural activity by the alpha rhythm in monkey visual cortex. *Current Biology*, *22*, 2313–2318. **DOI:** <https://doi.org/10.1016/j.cub.2012.10.020>, **PMID:** 23159599, **PMCID:** PMC3528834
- Spaak, E., de Lange, F. P., & Jensen, O. (2014). Local entrainment of alpha oscillations by visual stimuli causes cyclic modulation of perception. *Journal of Neuroscience*, *34*, 3536–3544. **DOI:** <https://doi.org/10.1523/JNEUROSCI.4385-13.2014>, **PMID:** 24599454, **PMCID:** PMC6608988

- Spelke, E., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, *99*, 605–632. **DOI:** <https://doi.org/10.1037/0033-295X.99.4.605>, **PMID:** 1454901
- Spratling, M. W. (2008). Reconciling predictive coding and biased competition models of cortical function. *Frontiers in Computational Neuroscience*, *2*, 1–8. **DOI:** <https://doi.org/10.3389/neuro.10.004.2008>, **PMID:** 18978957, **PMCID:** PMC2576514
- Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*, 745–756. **DOI:** <https://doi.org/10.1038/nrn3838>, **PMID:** 25315388
- Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, *13*, 403–409. **DOI:** <https://doi.org/10.1016/j.tics.2009.06.003>, **PMID:** 19716752
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M. M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, *11*, 1004–1006. **DOI:** <https://doi.org/10.1038/nn.2163>, **PMID:** 19160497, **PMCID:** PMC2747248
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press. <http://www.cs.ualberta.ca/sutton/book/ebook/the-book.html>, **DOI:** <https://doi.org/10.1109/TNN.1998.712192>
- Thomson, A. M. (2010). Neocortical layer 6: A review. *Frontiers in Neuroanatomy*, *4*, 13. **DOI:** <https://doi.org/10.3389/fnana.2010.00013>, **PMID:** 20556241, **PMCID:** PMC2885865
- Thomson, A. M., & Lamy, C. (2007). Functional maps of neocortical local circuitry. *Frontiers in Neuroscience*, *1*, 19–42. **DOI:** <https://doi.org/10.3389/neuro.01.1.1.002.2007>, **PMID:** 18982117, **PMCID:** PMC2518047
- Todorovic, A., van Ede, F., Maris, E., & de Lange, F. P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: An MEG study. *Journal of Neuroscience*, *31*, 9118–9123. **DOI:** <https://doi.org/10.1523/JNEUROSCI.1425-11.2011>, **PMID:** 21697363, **PMCID:** PMC6623501
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *The analysis of visual behavior* (pp. 549–586). Cambridge, MA: MIT Press.
- Urakubo, H., Honda, M., Froemke, R. C., & Kuroda, S. (2008). Requirement of an allosteric kinetics of NMDA receptors for spike timing-dependent plasticity. *Journal of Neuroscience*, *28*, 3310–3323. **DOI:** <https://doi.org/10.1523/JNEUROSCI.0303-08.2008>, **PMID:** 18367598, **PMCID:** PMC6670607
- Usrey, W. M., & Sherman, S. M. (2018). Corticofugal circuits: Communication lines from the cortex to the rest of the brain. *Journal of Comparative Neurology*, *527*, 640–650. **DOI:** <https://doi.org/10.1002/cne.24423>, **PMID:** 29524229, **PMCID:** PMC6131091
- Valpola, H. (2015). From neural PCA to deep unsupervised learning. In *Advances in independent component analysis and learning machines*. (pp. 143–171). Cambridge, MA: Academic Press.
- van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., et al. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *111*, 14332–14341. **DOI:** <https://doi.org/10.1073/pnas.1402773111>, **PMID:** 25205811, **PMCID:** PMC4210002
- VanRullen, R. (2016). Perceptual cycles. *Trends in Cognitive Sciences*, *20*, 723–735. **DOI:** <https://doi.org/10.1016/j.tics.2016.07.006>, **PMID:** 27567317
- VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences*, *7*, 207–213. **DOI:** [https://doi.org/10.1016/S0042-6989\(02\)00298-5](https://doi.org/10.1016/S0042-6989(02)00298-5), **PMID:** 12757822
- VanRullen, R., & Thorpe, S. J. (2002). Surfing a spike wave down the ventral stream. *Vision Research*, *42*, 2593–2615. **DOI:** [https://doi.org/10.1016/S0042-6989\(02\)00298-5](https://doi.org/10.1016/S0042-6989(02)00298-5), **PMID:** 12446033
- Varela, F. J., Toro, A., John, E. R., & Schwartz, E. L. (1981). Perceptual framing and cortical alpha rhythm. *Neuropsychologia*, *19*, 675–686. **DOI:** [https://doi.org/10.1016/0028-3932\(81\)90005-1](https://doi.org/10.1016/0028-3932(81)90005-1), **PMID:** 7312152
- Vinken, K., & Vogels, R. (2017). Adaptation can explain evidence for encoding of probabilistic information in macaque inferior temporal cortex. *Current Biology*, *27*, R1210–R1212. **DOI:** <https://doi.org/10.1016/j.cub.2017.09.018>, **PMID:** 29161556
- von Stein, A., Chiang, C., & König, P. (2000). Top-down processing mediated by interareal synchronization. *Proceedings of the National Academy of Sciences, U.S.A.*, *97*, 14748–14753. **DOI:** <https://doi.org/10.1073/pnas.97.26.14748>, **PMID:** 11121074, **PMCID:** PMC18990
- von Helmholtz, H. (1867/2013). *Treatise on physiological optics* (Vol. 3). North Chelmsford, MA: Courier Corporation.
- Waldert, S., Lemon, R. N., & Kraskov, A. (2013). Influence of spiking activity on cortical local field potentials. *Journal of Physiology*, *591*, 5291–5303. **DOI:** <https://doi.org/10.1113/jphysiol.2013.258228>, **PMID:** 23981719, **PMCID:** PMC3936368
- Walsh, K. S., McGovern, D. P., Clark, A., & O'Connell, R. G. (2020). Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the New York Academy of Sciences*, *1464*, 242–268. **DOI:** <https://doi.org/10.1111/nyas.14321>, **PMID:** 32147856, **PMCID:** PMC7187369
- Walter, W. G. (1953). *The living brain*. Oxford, UK: W. W. Norton.
- Watanabe, T., & Sasaki, Y. (2015). Perceptual learning: Toward a comprehensive theory. *Annual Review of Psychology*, *66*, 197–221. **DOI:** <https://doi.org/10.1146/annurev-psych-010814-015214>, **PMID:** 25251494, **PMCID:** PMC4286445
- Whittington, J. C. R., & Bogacz, R. (2019). Theories of error back-propagation in the brain. *Trends in Cognitive Sciences*, *23*, 235–250. **DOI:** <https://doi.org/10.1016/j.tics.2018.12.005>, **PMID:** 30704969, **PMCID:** PMC6382460
- Williams, R. J., & Zipser, D. (1992). Gradient-based learning algorithms for recurrent networks and their computational complexity. In Y. Chauvin & D. E. Rumelhart (Eds.), *Backpropagation: The theory, architectures and applications*. Hillsdale, NJ: Erlbaum.
- Wilson, J. R., Bose, N., Sherman, S. M., & Guillery, R. W. (1984). Fine structural morphology of identified X- and Y-cells in the cat's lateral geniculate nucleus. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, *221*, 411–436. **DOI:** <https://doi.org/10.1098/rspb.1984.0042>, **PMID:** 6146984
- Wimmer, R. D., Schmitt, L. I., Davidson, T. J., Nakajima, M., Deisseroth, K., & Halassa, M. M. (2015). Thalamic control of sensory selection in divided attention. *Nature*, *526*, 705–709. **DOI:** <https://doi.org/10.1038/nature15398>, **PMID:** 26503050, **PMCID:** PMC4626291
- Wiskott, L., & Sejnowski, T. J. (2002). Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, *14*, 715–770. **DOI:** <https://doi.org/10.1162/089976602317318938>, **PMID:** 11936959
- Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *Journal of Neuroscience*, *20*, RC63. **DOI:** <https://doi.org/10.1523/JNEUROSCI.20-06-j0002.2000>, **PMID:** 10704517, **PMCID:** PMC6772495

- Wurtz, R. H. (2008). Neuronal mechanisms of visual stability. *Vision Research*, *48*, 2070–2089. **DOI:** <https://doi.org/10.1016/j.visres.2008.03.021>, **PMID:** 18513781, **PMCID:** PMC2556215
- Xing, D., Yeh, C.-I., Burns, S., & Shapley, R. M. (2012). Laminar analysis of visually evoked activity in the primary visual cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *109*, 13871–13876. **DOI:** <https://doi.org/10.1073/pnas.1201478109>, **PMID:** 22872866, **PMCID:** PMC3427063
- Yu, C., & Smith, L. B. (2012). Embodied attention and word learning by toddlers. *Cognition*, *125*, 244–262. **DOI:** <https://doi.org/10.1016/j.cognition.2012.06.016>, **PMID:** 22878116, **PMCID:** PMC3829203
- Zhou, H., Schafer, R. J., & Desimone, R. (2016). Pulvinar-cortex interactions in vision and attention. *Neuron*, *89*, 209–220. **DOI:** <https://doi.org/10.1016/j.neuron.2015.11.034>, **PMID:** 26748092, **PMCID:** PMC4723640